

On The Ridiculousness of Notice and Consent: Contradictions in App Privacy Policies

Ehimare Okoyomon



Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2019-76

<http://www2.eecs.berkeley.edu/Pubs/TechRpts/2019/EECS-2019-76.html>

May 17, 2019

Copyright © 2019, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

On The Ridiculousness of Notice and Consent: Contradictions in App Privacy Policies

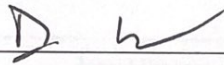
Ehimare Okoyomon

Research Project

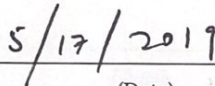
Submitted to the Department of Electrical Engineering and Computer Sciences,
University of California at Berkeley, in partial satisfaction of the requirements for
the degree of **Master of Science, Plan II.**

Approval for the Report and Comprehensive Examination:

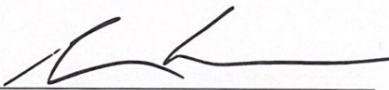
Committee



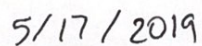
David Wagner
Research Advisor



(Date)



Serge Egelman
Second Reader



On The Ridiculousness of Notice and Consent: Contradictions in App Privacy Policies

Ehimare Okoyomon¹ Nikita Samarin¹ Primal Wijesekera^{1,2} Amit Elazari Bar On¹
Narseo Vallina-Rodriguez^{2,3} Irwin Reyes² Álvaro Feal^{3,4} Serge Egelman^{1,2}

¹University of California, Berkeley, ²International Computer Science Institute, ³IMDEA Networks Institute, ⁴Universidad Carlos III de Madrid

Abstract

The dominant privacy framework of the information age relies on notions of “notice and consent.” That is, service providers will disclose, often through privacy policies, their data collection practices, and users can then consent to their terms. However, it is unlikely that most users comprehend these disclosures, which is due in no small part to ambiguous, deceptive, and misleading statements. By comparing actual collection and sharing practices to disclosures in privacy policies, we demonstrate the scope of this problem.

Through analysis of 68,051 apps from the Google Play Store, their corresponding privacy policies, and observed data transmissions, we investigated the potential misrepresentations of apps in the Designed For Families (DFF) program, inconsistencies in disclosures regarding third-party data sharing, as well as contradictory disclosures about secure data transmissions. We find that of the 8,030 DFF apps (i.e., apps directed at children), 9.1% claim that their apps are not directed at children, while 30.6% claim to have no knowledge that the received data comes from children. In addition, we observe that 31.3% of the 22,856 apps that do not declare any third-party providers still share personal identifiers, while only 22.2% of the 68,051 apps analyzed explicitly name third parties. This ultimately makes it not only difficult, but in most cases impossible, for users to establish where their personal data is being processed. Furthermore, we find that 9,424 apps do not use TLS when transmitting personal identifiers, yet 28.4% of these apps claim to take measures to secure data transfer. Ultimately, these divergences between disclosures and actual app behaviors illustrate the ridiculousness of the notice and consent framework.

1 Introduction

Data protection and privacy regulations are largely informed by the Fair Information Practice Principles (FIPPs) – a set of

practices governing the collection and usage of personal information by different entities [37]. Central to FIPPs, and privacy regulations more generally, are the principles of *transparency* and *choice*, which are often presented as “notice and consent”, or informed consent. In the context of online privacy, service providers (such as websites or mobile applications) are often required to disclose their information collection practices to users and obtain their consent before collecting and sharing personal information. The most common mechanism of achieving this is by having the users consent to a privacy policy presented by the service provider.

Literature has demonstrated limitations of “notice and consent” [7, 10, 11, 48, 66] and recent regulations, such as the EU General Data Protection Regulation (GDPR) and California Consumer Privacy Act (CCPA), further require online services to provide users comprehensive information on data collection and sharing practices. Such information includes the type of personal information collected and shared, the purpose of the collection (in some cases), and the category of the recipient of the information [1, 20]. Such notice requirements play an important role in the mobile app ecosystem, where commonly used operating system permissions may inform users about potential data collection, but do not provide any insight as to who is the recipient, and for what purpose the information is collected.

However, due to the absence of stringent regulations, online services often draft privacy policies vaguely and obscurely, rendering the informed consent requirement ineffective. Moreover, related literature has demonstrated that privacy notices are often written at a college reading level, making them less comprehensible to average users [3, 29, 31, 63]. Even if privacy policies were more comprehensible, prior work suggested that users would still need to spend over 200 hours per year on average reading every privacy policy associated with every website they visit [49]. It comes as no surprise, therefore, that in reality, very few users choose to read privacy notices in the first place [50, 52, 55].

Our work aims to further demonstrate the inadequacy of privacy policies as a mechanism of notice and consent, fo-

A version of this work was published at the 2019 IEEE Workshop on Technology and Consumer Protection (ConPro '19).

ocusing on Android smartphone applications (‘apps’). Literature has shown questionable privacy behaviors and collection practices across the mobile app ecosystem [61]. This paper explores whether such specific questionable collection practices are represented in the privacy policies and disclosed to users. While past work has focused separately on app behavior analysis at practice [5, 17, 22, 53, 60, 61, 64, 73] or analysis of privacy policies [9, 30, 76, 77, 81], we aim to bridge this gap by considering these two problems in tandem. In other words, we compliment the dynamic analysis results, focusing on what is collected and with whom it is shared, with an analysis of whether users were adequately informed about such collection.

In this paper we focus on three classes of discrepancies between collection at practice (de facto) and as per the online service’s notice (de jure):

- First, we examine mobile apps that participate in the Google Play Store’s ‘Designed for Families’ (DFF) program and regulated under the Children Online Privacy Protection Act (COPPA), meaning their target audience includes children under the age of 13 [69]. We find that a substantial number of apps *targeted at children* include clauses in their privacy policy either claiming to not have knowledge of children in their audience, or outright prohibitions against the use of their apps by children.
- The second aspect we are interested in analyzing is the disclosure of third-party services that receive and process user information. Regulations like GDPR (Article 13 1.e) and CCPA require developers to explicitly notify users about the recipients of information, either their names or categories. We explored how many app developers include information about their third-party affiliates in the privacy policy and how many of them explicitly name them.
- Third, privacy policies often represent to users they implement reasonable security measures. At a minimum, one such measure should include TLS encryption. Protecting users’ data using reasonable security measures is a regulatory requirement under COPPA, CCPA, and GDPR (Article 32). We explored how many apps potentially fail to adhere to their own represented policies, by transmitting data without using TLS.

2 Background and Related Work

2.1 Children’s Privacy

In the Children’s Privacy sections of this paper, we analyze the transparency of mobile applications published in Google Play’s Designed for Families Program. Particularly, we focus on identifying potential inconsistencies between DFF applications’ runtime behavior and the information publicly disclosed

in their privacy policies. We provide an overview of Google Play Store’s DFF program and discuss below the most related work in the area of privacy violations in children apps, third-party libraries and automatic analysis of privacy policies.

Designed For Families (DFF): The DFF program is a platform for developers to present age appropriate content for the whole family, specifically users below the age of 13. By going through an optional review process, applications submitted to this program declare themselves to be appropriate for children and are able to list themselves under such family-friendly categories. Although the *primary* target audience of these apps may not be children, according to [69], “if your service targets children as one of its audiences - even if children are not the primary audience - then your service is ‘directed to children.’” As a result, DFF applications are all directed at children and thus required to be “compliant with COPPA (Children’s Online Privacy Protection Rule), the EU General Data Protection Regulation 2016/679, and other relevant statutes, including any APIs that your app uses to provide the service” [24].

Privacy Violations in Children’s Apps: There have been some efforts by the research community to investigate children’s privacy in social media [46], smart toys [72, 78] and child-oriented mobile apps [61]. Furthermore, researchers have studied how well developers that target children as their audience comply with specific legislation such as COPPA [61] and the FTC has already established cases against numerous app developers gathering children’s data [68] [70] [12].

2.2 Third Parties and Tracking

The Mobile Advertising and Tracking Ecosystem: It is well known that mobile apps rely on third parties to provide services such as crash reports, usage statistics or targeted advertisement [59]. In order to study such ecosystems, previous work has relied on analysis of mobile traffic from several vantage points such as ISP logs, heavily instrumented phones or devices that use VPN capabilities to observe traffic [60, 64, 73]. Another way of studying the presence of such libraries is by performing dynamic analysis (i.e., studying runtime behavior of the app) or static analysis (i.e., studying application code to infer actual behavior). Both these approaches have long been used to analyze privacy in mobile apps [5, 17, 22, 53].

Data Over-Collection: With the prevailing existence of data collection and third party data sharing, there comes the worry of data over-collection in which smartphone apps collect more users’ data than needed to perform its original function. Related research projects [17, 41, 74] have conducted experiments and studies of frequent data over-collection cases and have even presented approaches to eradicate data over-collection but the problem still exists today. While third party services may add great value, this comes at an enormous privacy cost. As a result, there have been, and continue to be, policy debates about the regulation of third-party tracking [19, 32, 40, 67].

2.3 Secure Transmission

TLS has become somewhat of a standardized method of secure data transmissions nowadays. Although prior work has shown that both using default OS APIs and third party libraries for TLS implementation have their weaknesses [58], TLS has been easily integrated into the toolkit of app developers. Consequently, TLS usage has seen an increase in adoption in recent years and its usage and integration in popular platforms has been growing [13,25] and there have even been several efforts in further securing TLS by studying its usability and implementation, and even outlining its transition in these categories over the years [35,39]. However, even with these toolkits being widely available, we still see prevalent issues of applications transmitting information, many times sensitive or personally-identifying, without using this secure transmission method, thus threatening the privacy of their users [18,61]. Despite the existing legislature in place [20,69], this makes one wonder why there is not more stringent regulation on applications to deploy these secure transmission methods.

2.4 App Analysis

The two main approaches in understanding the behaviour of software applications can be broken up into two categories: *static* and *dynamic* analysis. Static analysis refers to studying the code and all its possible flows in order to uncover possible behaviours. For example, numerous research studies have utilized static analysis in search of suspicious and potentially malicious application behaviors [5,6,28,34,80]. However, the issue with static analysis is that an application’s runtime behaviour may differ greatly from its hypothetical behaviours. In doing this analysis, one cannot always predict every use case and it is possible that is even possible that the structure of the code changes, rendering this method analysis vulnerable to different attacks [14,21,47]. Fortunately, studies have also been known to use dynamic analysis, which has the ability to execute applications and monitor their behaviours at runtime [43,44,47,75,78]. However, this dynamic analysis is much harder to do at scale as some apps are adverse to running in virtual environments while others may need real user input. Nonetheless, there still exist tools to circumvent the majority of these problems [26,38,56,57]. Our data in this paper was obtained by previous work that employed dynamic analysis as their primary means of app analysis and information retrieval [4].

2.5 Privacy Policy Analysis

Automatic Analysis of Privacy Policies: The research community has previously studied publicly available privacy policies proving that they are usually written in a “*legal language*” that makes it difficult for average users to understand them [33,54,65]. Researchers have made positive efforts to

extract relevant information and features from privacy policies using crowd-sourcing campaigns [9,77], Natural Language Processing (NLP) [81] and deep-learning [30,76] techniques to present this information in more accessible ways to users [8,15,30,42,51,79]. While we are aware these solutions exist, they prove to not be helpful in all cases as shown in our example below. Therefore, despite these positive efforts, privacy policies remain a confusing source of information for average users and these documents are even unavailable in many webpages and mobile apps [16,30,81].

Polisis: One such service that enables automatic analysis of privacy policies is Polisis [30]. Polisis enables users to submit websites for review, which it then crawls to find the associated privacy policies. It then analyzes the texts of these policies and produces a summary of data collection practices of the submitted websites. We initially considered using Polisis to parse the privacy policies and extract information relevant to our current work. However, we encountered an issue with rate-limiting as we created too many requests and were temporarily banned from using the service. Moreover, we discovered that the performance of Polisis on establishing the sections related to children’s privacy was suboptimal. Out of a random subset of 100 policies that we categorized as claiming not being directed at children, Polisis was able to identify only 22% of policies. Therefore, due to the overhead in the data generation using Polisis as well as poor performance on tasks relevant to our work, we decided to perform our own policy analysis.

Our paper utilizes previous works and ideologies from privacy policy and app analysis in a combined effort in order to shine light on, and add depth to, these issues in Children’s Privacy, Third-Party Partnerships, and Secure Communications. Ultimately, this multi-faceted approach provides a better understanding of the effect of the privacy framework of notice and consent.

3 Dataset and Methodology

In this section, we describe the datasets and methods that we use to establish the misrepresentations occurring in privacy policies. We use a publicly available app analysis dataset [4] to evaluate the types of data that mobile apps access and share with other parties over the Internet. We then explain how we leverage this dataset to extract meaning and trends from the apps’ behaviors and privacy declarations.

3.1 AppCensus Dataset

In our work, we rely on the AppCensus dataset available at [4]. AppCensus is a tool that analyzes Android apps from the Google Play Store in order to identify the personal information that apps access and share with other parties over the

Table 1: Number of observed apps for different types of analysis.

Description	Observed App #	Sample Size
Participate in DFF program	8,030	68,051
Claim not to target children	728	8,030
Claim no knowledge of children data	2,457	8,030
Mention third parties	45,195	68,051
Provide names of third parties	15,106	45,195
Undisclosed sharing (third parties not mentioned)	7,147	22,856
No TLS usage in data transmission	9,424	68,051
Claim to secure data transmission	2,680	9,424

Internet. It leverages dynamic analysis techniques to automatically analyze an application’s runtime and network behavior. AppCensus runs in a highly instrumented Android operating system that is able to detect the personal information linked to the test phone. Then, the tool runs each app in the test environment and inspects the network communications using a monitoring tool that records the source, destination, and content of network flows. Thus, we have access to the network flows of every app and we can find whether it shares personal data over the network and with whom. AppCensus also fetches and stores privacy notices of each analyzed app, which we use to identify possible mismatches between the stated and actual app behavior.

AppCensus stores analysis results and other metadata in a large-scale dataset that contains detailed information about the runtime behavior and traffic patterns of Android apps. As of January 2019, it includes information about 68,051 apps published on the Google Play Store. We compare the information provided in the privacy policies of these apps with their actual data sharing practices, as described in the next subsection.

3.2 Policy Analysis

In our project, we focus on three types of misrepresentations that occur in privacy notices of mobile apps. Table 1 shows the total number of apps from the AppCensus dataset that we examine and the number of observations that we obtain for different types of analysis. For misrepresentations concerning children’s privacy, we analyze 8,030 apps participating in Google’s DFF program out of all 68,051 available apps. For third-party sharing practices and for TLS usage, we use the entire dataset of 68,051 apps.

We analyze the text of privacy policies to identify potential misrepresentations. To verify compliance with COPPA, we first narrow our search to only include apps in Google’s DFF category with any combination of the keywords "child", "kid", "COPPA", and "minor." Next, we manually read and process the policies for a subset of 200 DFF apps, focusing primarily on these keywords and frequently-used phrases and expressions. This allows us to identify commonalities and phrases

that could generally be categorized as:

1. Those that make no distinction between children under 13 and other users
2. Those that clearly indicate that the application either
 - (a) is not directed at children; or
 - (b) does not knowingly collect personally identifiable information from children under 13
3. Indeterminate (i.e. phrases that need further manual processing, perhaps due to confusing language)

Enumerating these phrases into categories allows us to create complex strings representative of these groups and further categorize other policies, most commonly those of type 2 and 3, by searching for other policies with matching substrings. This approach enables us to classify over 400 more policies. From the remaining list, we are able to loosen the specificity of the search strings by simply searching for word combinations such as "not knowingly", "not targeted to/at", "do not address" and their variations to categorize policies containing these phrases in lists. We further process these initial lists to remove irrelevant policies if such phrases were not made in the context of children’s privacy. Fortunately, this method produces quite a large throughput of roughly 700 more apps. The remaining policies can be read manually and largely represent policies in category 1. Furthermore, after obtaining and classifying the policy data, we look to draw conclusions directly from the classifications and text, as well as by cross-checking policy declarations with observed network transmissions collected from the respective applications.

We are further interested in exploring how many app developers disclose their information sharing practices. We look at all 68,051 available apps, aiming to collect the relevant clauses on information sharing with third-party services from their privacy policies and to determine whether the names (as opposed to categories) of those third-party recipients of information are disclosed.

First, we analyze the texts of privacy policies using regular expressions. In particular, we are interested to see whether

any part of the text matches the phrase "third parties" or any variation thereof (e.g. "affiliate" or "partner" instead of "third party"). While this approach is less sophisticated than techniques used by other automatic policy analyzers [2, 9, 30, 76, 77, 81], we believe that it is sufficient to identify clauses that contain any information about third parties. Focusing on the sentences matched by our regular expressions, we aim to identify which, if any, affiliates are explicitly named by the app developer. This requires solving the *named-entity recognition* (NER) problem and while state-of-the-art NER systems produce near human-level performance [23, 36, 45, 62], we discovered that they are ill-suited for reliably recognizing the names of analytics and advertising network companies.

Instead, we use the app analysis dataset (68,051) to determine all domains that receive data from mobile apps. This produces a list of 9,672 domains, including known analytics and advertising networks, such as *crashlytics.com*, *vungle.com*, *flurry.com*, *mopub.com* among many others. By separating the domain names and performing a manual review, we end up with a list of 7,826 domains. The entries in this list are matched against the text in the privacy policies to determine the third-party service providers named by the app developer.

Finally, we want to ensure that app developers comply with their own policies whenever they promise to take reasonable steps to secure user data from unauthorized access. ‘Reasonable security measures’ is a broad concept and includes different techniques, such as data encryption, regular security patches, access control, among many others. We focus on secure data sharing, which we believe belongs to ‘reasonable security measures’, as we have access to data transmission information from the AppCensus dataset.

In order to achieve this, we first identify mobile apps that transmit personal information over the Internet without using TLS. We then analyze their privacy policies, identifying parts of the text that mention personal data. This is again done using regular expressions, matching "personal information", "personally identifiable information" and variants thereof. Finally, sentences containing information about personal data are scanned for specific key phrases (e.g. "security measures", "unauthorized disclosure", "reasonable steps to secure", "transmission", etc.), that provide security guarantees concerning data transmission.

4 Results

We report our analysis along three aforementioned dimensions: Children’s Privacy, Third-party Service Providers, and Secure Data Transmission.

4.1 Children’s Privacy

For the Children’s Data Privacy analysis, we looked at 8,030 apps in the Designed For Families program. Out of these apps, we found that there are 728 apps (9.1%) that claim they are not targeted at children and 2,457 (30.6%) that claim no knowledge of collecting any data from children under 13, with some overlap in apps that do both. In fact, only 4,649 (57.9%) mention any combination of the keywords "child", "kid", "coppa", and "minor". Within this group there are even applications such as "Cinderella Classic Tale Lite" and "Dino Puzzle - free Jigsaw puzzle game for Kids" that make no commentary on children’s use directly in their privacy policies but instead simply contain one or more of these keywords either in the name of the app or in a header on the website. Thus, it is interesting to note the sheer number of apps designed for children that do not even mention kids or children in their policies.

From these 2,457 apps in DFF that claim no knowledge of collecting data from children, we observed 68,594 network transmissions. Out of 2,457 apps, 1,728 (70.3%) of them transfer data types such as the android advertising ID (AAID), IMEI, geolocation, and WiFi mac address. Figure 1 displays the breakdown of the number of apps that sent each data type as well as the most ubiquitously sent identifiers.

Since these apps are present in the DFF program, they are likely all covered by COPPA and thus are responsible for catering to an audience that includes minors. However, it is both confusing and troubling to see their policies contradict what they have acknowledged under DFF.

One such of these apps is “Smart Games for Kids for Free” made by the developer DEVGAME KIDS. Despite the very obvious advertising directed at children as inferred by the application icon and name, in addition to its declaration in the Google Play Store being for Ages 8 & Under, DEVGAME KIDS’ privacy policy claims they do not knowingly allow such persons to access their Services. In addition to this, they claim to not knowingly collect or solicit personal information from children, but we have observed them transmitting the AAID, androidid, and geolaton datatypes.



Smart Games for Kids for Free
 DEVGAME KIDS Educational Education ★★★★★ 4,971
 Everyone Ages 8 & Under
 Contains Ads - Offers In-app purchases
 You don't have any devices. Add to Wishlist Install

Through these Services, DEVGAME does not knowingly collect or solicit personal information from anyone under the age of 13 or knowingly allow such persons to access those Services.

Datatypes Transmitted:	AAID, androidid, geolaton
-------------------------------	----------------------------------

We further analyze this case amongst others in Discussion.

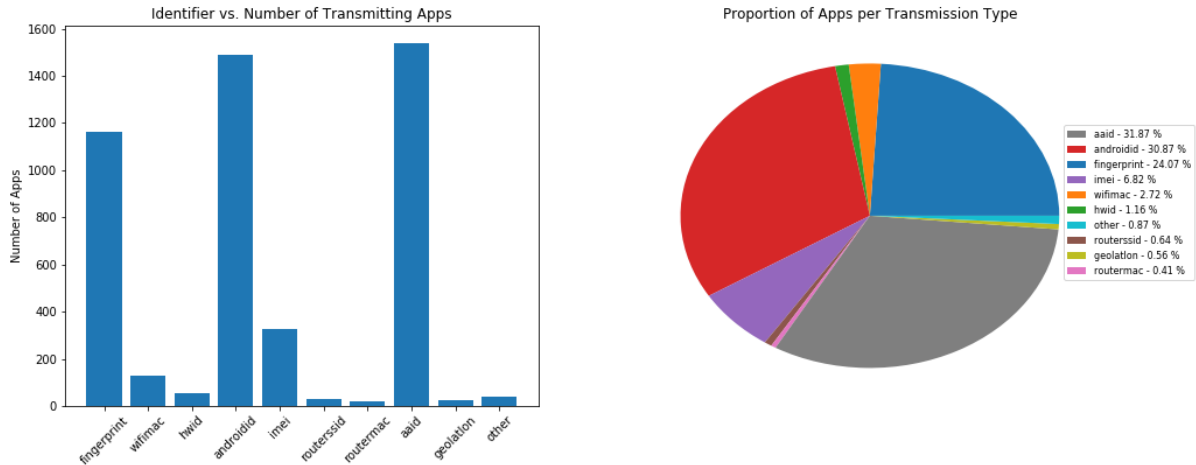


Figure 1: The number of apps that transmitted each datatype and the proportion of datatypes sent by transmitting apps, respectively.

4.2 Third-party Service Providers

We also identify apps that do not reveal the names of affiliated third parties in their privacy policies. We start by locating apps that mention third-party service providers. From there, we narrow this list only to include apps that explicitly name at least one third-party partner.

In our corpus, 45,195 (66.4%) mention third-party affiliates, which suggest that the remaining 22,856 apps should not transmit any personal data to outside domains. However, out of these 22,856 apps, 7,147 (31.3%) of them still share user identifiers with other service providers without giving notice to the users. In addition, we discover that only 15,106 apps (22.2% of 68,051) explicitly name their third-party affiliates. For instance, the game development company Kwalee describes its analytics service providers in the following way:

We use third party providers Fabric (for crash analytics), AppsFlyer and Tenjin (user attribution and analytics), and SOOMLA (for analytics). Fabric, AppsFlyer and SOOMLA may collect analytics on our behalf and in accordance with our instructions, and their applicable privacy notices.

One of the Kwalee’s app named Looper! does share Android Advertising ID (AAID) with the aforementioned third parties. However, we also discover that it also sends Android ID to the same parties, which is a unique identifier that could be used to track users over time and across apps, and can only be reset by performing a factory reset of their phone. Furthermore, Looper! transmits data to two undisclosed third-party affiliates, including SupersonicAds (advertisement platform) and Kochava (analytics platform).

Alternatively, in regards to third-party affiliates, the company Azar does not go so far as to mention the identities of

their affiliates as they simply state:

In certain instances, Hyperconnect, Inc. and Azar may have contracts with third parties to provide products and/or services in connection with the Service. In those cases, Azar may share your personal information with Azar’s third-party service providers, but only in connection with providing the services to you.

According to the measurements obtained from AppCensus, Azar collects and transmits AAID and Android ID. This information is sent to Facebook and Branch, a mobile measurement and deep linking platform.

4.3 Secure Data Transmission

Using the AppCensus dataset, we discover that 36,107 apps that are available on Google Play Store transmit personal data over the network. As of January 2019, 9,424 of these apps (26.1%) do not use TLS when transmitting personal identifiers. Although this fact is alarming in itself, we also investigate whether the developers of these apps make any deceptive claims in their privacy policies.

Out of those 9,424 apps, 2,680 apps claim to take measures to secure data transmission, but fail to employ TLS when transmitting PII. For instance, the game developer Animoca Brands describes their procedures in the following way:

We understand the importance of keeping your information secure and safe. We have implemented commercially reasonable security measures to secure and safeguard your information and to prevent the loss, misuse, destruction and alteration of information under our control.

However, we observe that one of the games (Astro Boy Siege: Alien Attack) transmits the AAID, Android ID and the IMEI number without encrypting data in-transit to a number of IP addresses, including one belonging to Unity. This app also transmits the same data using TLS encryption to other third-party service providers, such as Chartboost, Adzcore, Flurry, Muneris, Tapjoy, among others.

5 Discussion

We looked at privacy policies for any contradictions from their own behavior and for important information missing in the policies. On a high level, we found that developers a) contradict themselves between what they mention in the policy and what they acknowledge in the Google Play Store, b) are not comprehensive in what they are claiming about their data sharing practices, and c) do claim to secure communication when they actually do not.

5.1 Children’s Privacy

The most troublesome finding from looking at the privacy policies of apps in DFF is that their own policies do not acknowledge their audience as containing kids. We found 9.1% of apps under DFF contain phrases in their policies that indicate their developers do not mean for these products to be used by children, despite their developers proactively targeting children.

*"We do not specifically market to children under 13"
"These Services do not address anyone under the age of 13"*

Although some phrasings are ambiguous about developers’ intent to exclude children from their products, we believe confusing language still represents a form of misrepresentation. Our analysis reveals a set of policy misrepresentations as a result of unclear—and at times contradictory—language. Many privacy policies already pose difficulties for consumers to comprehend, but those challenges are further exacerbated when policies send mixed messages or are wrong. Such an effect runs counter to what privacy policies are supposed to achieve in increasing transparency. For instance, the privacy policy from the developer DEVGAME KIDS claims:

Most of DEVGAME’s Services are for kids audience and are directed to children under the age of 13. Through these Services, DEVGAME does not knowingly collect or solicit personal information from anyone under the age of 13 or knowingly allow such persons to access those Services.

A consumer, such as a concerned parent, reading this policy would be left with conflicting impressions about whether DEVGAME intends for children to use their apps. Such a confus-

ing and contradictory policy might indicate the level of importance (or lack thereof) this developer places on user privacy.

A particular set of contradictions occur in the 30.6% of apps under DFF whose policies claim that they do "not knowingly" collect information. This is concerning because either this is done as a purposefully deceptive practice or app developers are truly unaware of what data their apps collect. Nonetheless, both explanations prove problematic when trying to conform to privacy laws such as COPPA and GDPR.

However, barring this, there also exists a variety of other examples of misleading and misrepresenting behaviors of these applications, as viewed from their policies and network transmissions. For example, the privacy policy for the app "Pony Crafting - Unicorn World" claims they do not knowingly collect information from children. However, they later mention that they can collect both personal and non-personal information, and that they do not actually know the source of their data:

...this information will be anonymous and we will not be able to tie this information to a specific user; this is a byproduct of an anonymous system not allowing us to distinguish between those over 13 and those under 13

In order to confirm if the transmission of children’s information really occurs, we searched the database for network traffic observed from applications whose policies do not "knowingly" collect from children. In doing this, we found 68,594 distinct transmissions across the 1,728 DFF apps that sent device identifiers across the network, including 1,979 transmissions from 10 apps that share the same policy.¹ Among these datatypes, we found identifiers such as the AAID, IMEI, geolocation, and WiFi mac address. While it is easy to recognize why the sharing of geolocation could be problematic, the issue with other identifiers is a little more subtle. For example, looking at the AAID, the most transmitted datatype observed as recorded in Figure 1, Google’s terms of use actually indicates that "the advertising identifier must not be connected to personally-identifiable information or associated with any persistent device identifier (for example: SSAID, MAC address, IMEI, etc.) without explicit consent of the user" [27]. However, 271 apps were observed sending both AAID and IMEI, while 120 apps sent both AAID and MAC address. In addition, under the 2013 amendments to the COPPA rule, persistent identifiers such as the IMEI and MAC addresses are considered personally identifying if they "can be used to recognize a user over time and across different websites or online services" [71]. Furthermore, there are certain rules that developers and third party services must follow when legitimately collecting personally identifying information under COPPA, and by declaring they do not knowingly collect data from

¹<http://www.vascogames.com/vasco-games-privacy-policy/> as accessed on 2018/11/25

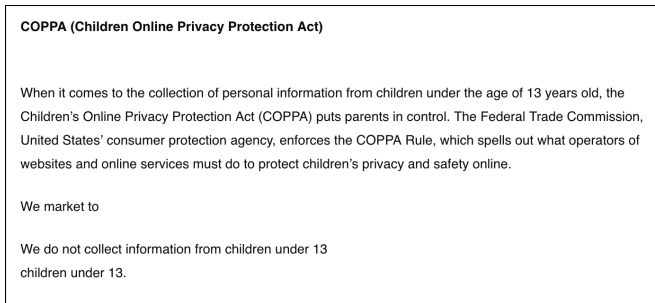


Figure 2: The Copy-and-Paste formatting error in the privacy policies of the app “Color Clues” and 30+ others analyzed in DFF.

children, these apps are implicitly admitting to not following these guidelines.

Since these results all come from automated runs of a UI Monkey that creates pseudo-random swipes and gestures, our testing employed no sophisticated methods to overcome age-gating, parental consent, or any other barriers. Therefore, we believe this behavior can be akin to that of a child and thus the resulting transmissions would be expected to occur with real users. As a result, we believe the claim of “not knowingly” collecting information is misleading because the applications do very little to verify that collection does not come from children despite explicitly marketing them to children. The majority of the applications also detail that if data from minors is found or reported to them they will delete the information and block the associated account, but we strongly feel this is impractical from a usability standpoint as the majority of users will not actually know what information is being sent in the background and thus cannot make this report.

Furthermore, another transient detail we observed during analysis, as briefly mentioned above in the Methodology, is the presence of repeated sentences across policies. In fact, numerous times many of the *exact same* phrases are used, which becomes particularly obvious when they include the same spelling, grammar, or formatting mistakes such as in the second part of the privacy policy of the app “Color Clues”, and at least 30 other applications, displayed in Figure 2.

Although it makes our analysis more efficient, the only highly probably cause for this behavior is that companies and developers are not actually creating their own privacy policies but instead just copy and pasting sections from others found online. Regardless, this observation shows the low level of priority these companies place on making privacy policies for the protection of their users. As a result, this begs the questions of whether application developers are even aware of the statements in their policies (such as in the DEVGAME example above), as well as if they are even aware of the laws they must abide by. They could perhaps be just including the minimal information necessary, as determined by examining

another application’s policy, in order to escape under the radar. This hypothesis is further supported by our findings when we look at the application “Cami’s Dots” that, instead of creating their own privacy policy, decided it was sufficient to provide a link pointing to github.com’s privacy policy page.² Thus, a potential future work stemming from these findings could be to investigate the percentage of applications that are plagiarised and the degree to which this occurs.

5.2 Third Parties

Our analysis has demonstrated that mobile apps do not provide sufficient information to users about their third-party partners. Although around 66% of mobile apps mention third-party data recipients, only around 22% actually disclose the names of those recipients. Out of remaining 22,856 apps, 31.3% share user data with third parties without providing notice to user. This is concerning, as the users are unable to learn about the information protection practices of the services that are provided access to their personal data. This raises questions about the feasibility of ensuring compliance with data protection legislation, as any potential contradictions within the policy of the app and that of the third party would go unnoticed.

Another common theme is the desire of app developers to relieve themselves of any responsibility for ways in which personal information about the users is handled once it leaves the application. In particular, one of such examples is in the following clause taken from the privacy policy of the app “Drum Pads 24 - Music Maker”:

We are not responsible or liable for the privacy policies or practices of third parties.

However, we believe that app developers should still be accountable for the actions performed by third parties, to which they send users’ personal information. It is possible that the desire to avoid liability is what motivates app developers not to provide names of their partners in the first place.

5.3 Secure Transmission

As mentioned previously, 26.1% of 36,107 mobile apps that send personal data do not use TLS when transmitting user identifiers and of these, 28.4% claim to take reasonable security measures. The reasons for making such claim might include negligence (e.g. not knowing that TLS is not configured, or copy-pasting the policy from another app) or malicious intention (e.g. to create a semblance of security without using proper means of protection). Most of these apps also include a clause explaining how it is impossible to completely ensure the security of data transmission:

²<https://help.github.com/articles/github-privacy-statement/> as accessed on 2018/11/25. Since 2018/12/18, Cami’s Dots’s policy moved to <https://github.com/LTPProjects/Cami-s-Dots/compare/master...jigglytep:patch-1>

But remember that no method of transmission over the internet, or method of electronic storage is 100% secure and reliable, and we cannot guarantee its absolute security.

As in the third party case, this clause also demonstrates the desire for app developers to avoid responsibility for negligent data protection standards. While it is true that it is impossible to be secure against all types of attacks on user data, we believe that using TLS for data transmission is one of the most basic steps that should be adopted by all app developers.

6 Conclusion

This paper accentuates the degree in which the privacy framework of notice and consent is flawed by analyzing Google Play Store apps and comparing their privacy policies with their behavior. Our analysis specifically focuses on highlighting the misrepresentation and lack of information that exists in of apps in the Designed for Families program, apps that interact with third parties, as well as apps that claim to utilize secure data transmission precautions, ultimately showing the level of carelessness and lack of priority when it comes to protecting consumer privacy.

Acknowledgments

This work was supported by the U.S. National Security Agency’s Science of Security program (contract H98230-18-D-0006), the Department of Homeland Security (contract FA8750-18-2-0096), the National Science Foundation (grants CNS-1817248 and CNS-1564329), the European Union’s Horizon 2020 Innovation Action program (grant Agreement No. 786741, SMOOTH Project), the Rose Foundation, the Data Transparency Lab, and the Center for Long-Term Cybersecurity at U.C. Berkeley.

References

- [1] Assembly Bill No. 375, Chapter 55. *California Legislative Information*, 2018.
- [2] Waleed Ammar, Shomir Wilson, Norman Sadeh, and Noah A Smith. Automatic categorization of privacy policies: A pilot study. *School of Computer Science, Language Technology Institute, Technical Report CMU-LTI-12-019*, 2012.
- [3] Annie I Anton, Julia Brande Earp, Qingfeng He, William Stufflebeam, Davide Bolchini, and Carlos Jensen. Financial privacy policies and the need for standardization. *IEEE Security & privacy*, 2(2):36–45, 2004.
- [4] AppCensus AppSearch. <https://search.appcensus.io/>. Accessed: 2019-03-26.
- [5] Steven Arzt, Siegfried Rasthofer, Christian Fritz, Eric Bodden, Alexandre Bartel, Jacques Klein, Yves Le Traon, Damien Octeau, and Patrick McDaniel. Flowdroid: Precise context, flow, field, object-sensitive and lifecycle-aware taint analysis for android apps. *Acm Sigplan Notices*, 49(6):259–269, 2014.
- [6] Vitalii Avdiienko, Konstantin Kuznetsov, Alessandra Gorla, Andreas Zeller, Steven Arzt, Siegfried Rasthofer, and Eric Bodden. Mining apps for abnormal usage of sensitive data. In *Proceedings of the 37th International Conference on Software Engineering-Volume 1*, pages 426–436. IEEE Press, 2015.
- [7] Solon Barocas and Helen Nissenbaum. On notice: The trouble with notice and consent. 2009.
- [8] Guido Boella, Luigi Di Caro, Michele Graziadei, Loredana Cupi, Carlo Emilio Salaroglio, Llio Humphreys, Hristo Konstantinov, Kornel Marko, Livio Robaldo, Claudio Ruffini, et al. Linking legal open data: breaking the accessibility and language barrier in european legislation and case law. In *Proceedings of the 15th International Conference on Artificial Intelligence and Law*, pages 171–175. ACM, 2015.
- [9] Travis D Breaux and Florian Schaub. Scaling requirements extraction to the crowd: Experiments with privacy policies. In *Requirements Engineering Conference (RE), 2014 IEEE 22nd International*, pages 163–172. IEEE, 2014.
- [10] Fred H Cate. The failure of fair information practice principles. 2006.
- [11] Fred H Cate. The limits of notice and choice. *IEEE Security & Privacy*, 8(2):59–62, 2010.
- [12] U.S. Federal Trade Commission. Two App Developers Settle FTC Charges They Violated Children’s Online Privacy Protection Act. <https://www.ftc.gov/news-events/press-releases/2015/12/two-app-developers-settle-ftc-charges-they-violated-c> 2016. Accessed: September 26, 2017.
- [13] K Conger. Apple will require https connections for ios apps by the end of 2016. <https://techcrunch.com/2016/06/14/apple-will-require-https-connections-for-ios-apps-by-2016>. Accessed: May 16, 2019.
- [14] Andrea Continella, Yanick Fratantonio, Martina Lindorfer, Alessandro Puccetti, Ali Zand, Christopher Kruegel, and Giovanni Vigna. Obfuscation-resilient privacy leak detection for mobile apps through differential analysis. In *NDSS*, 2017.

- [15] Michael Curtotti and Eric McCreath. A right to access implies a right to know: An open online platform for research on the readability of law. *J. Open Access L.*, 1:1, 2013.
- [16] Martin Degeling, Christine Utz, Christopher Lentzsch, Henry Hosseini, Florian Schaub, and Thorsten Holz. We value your privacy... now take some cookies: Measuring the gdpr's impact on web privacy. *arXiv preprint arXiv:1808.05096*, 2018.
- [17] William Enck, Peter Gilbert, Seungyeop Han, Vasant Tendulkar, Byung-Gon Chun, Landon P Cox, Jaeyeon Jung, Patrick McDaniel, and Anmol N Sheth. Taint-droid: an information-flow tracking system for realtime privacy monitoring on smartphones. *ACM Transactions on Computer Systems (TOCS)*, 32(2):5, 2014.
- [18] BENJAMIN ERIKSSON and JONAS GROTH. On the road with third-party apps.
- [19] ComisiÃ³n Europea. Commission proposes a comprehensive reform of the data protection rules. *Data Protection*, 2012.
- [20] Regulation (EU) 2016/679 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). *Official Journal of the European Union*, L119:1–88, May 2016.
- [21] Parvez Faruki, Ammar Bharmal, Vijay Laxmi, Manoj Singh Gaur, Mauro Conti, and Muttukrishnan Rajarajan. Evaluation of android anti-malware techniques against dalvik bytecode obfuscation. In *2014 IEEE 13th International Conference on Trust, Security and Privacy in Computing and Communications*, pages 414–421. IEEE, 2014.
- [22] Yu Feng, Saswat Anand, Isil Dillig, and Alex Aiken. Appscopy: Semantics-based detection of android malware through static analysis. In *Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering*, pages 576–587. ACM, 2014.
- [23] Jenny Rose Finkel, Trond Grenager, and Christopher Manning. Incorporating non-local information into information extraction systems by gibbs sampling. In *Proceedings of the 43rd annual meeting on association for computational linguistics*, pages 363–370. Association for Computational Linguistics, 2005.
- [24] Google. Developer policy center. [Online; accessed 2019-03-26].
- [25] Google 2019. Https encryption on the web. <https://www.google.com/transparencyreport/https/>, 2019. Accessed: May 16, 2019.
- [26] Google, Inc. UI/Application Exerciser Monkey. <https://developer.android.com/tools/help/monkey.html>.
- [27] Google, Inc. Usage of android advertising id. <https://play.google.com/about/monetization-ads/ads/ad-id/>, 2017. Accessed: November 30, 2017.
- [28] Michael I Gordon, Deokhwan Kim, Jeff H Perkins, Limei Gilham, Nguyen Nguyen, and Martin C Rinard. Information flow analysis of android applications in droidsafe. In *NDSS*, volume 15, page 110, 2015.
- [29] Mark A Graber, Donna M D'alessandro, and Jill Johnson-West. Reading level of privacy policies on internet health web sites.(brief report). *Journal of Family Practice*, 51(7):642–646, 2002.
- [30] Hamza Harkous, Kassem Fawaz, Rémi Leuret, Florian Schaub, Kang G. Shin, and Karl Aberer. Polisis: Automated analysis and presentation of privacy policies using deep learning. In *27th USENIX Security Symposium (USENIX Security 18)*, pages 531–548, Baltimore, MD, 2018. USENIX Association.
- [31] Mark Hochhauser. Lost in the fine print: Readability of financial privacy notices. *Retrieved November, 27:2009*, 2001.
- [32] White House. Consumer data privacy in a networked world: A framework for protecting privacy and promoting innovation in the global digital economy. *White House, Washington, DC*, pages 1–62, 2012.
- [33] Carlos Jensen and Colin Potts. Privacy policies as decision-making tools: an evaluation of online privacy notices. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, pages 471–478. ACM, 2004.
- [34] Jinyung Kim, Yongho Yoon, Kwangkeun Yi, Junbum Shin, and SWRD Center. Scandal: Static analyzer for detecting privacy leaks in android applications. *MoST*, 12:110, 2012.
- [35] Platon Kotzias, Abbas Razaghpanah, Johanna Amann, Kenneth G Paterson, Narseo Vallina-Rodriguez, and Juan Caballero. Coming of age: A longitudinal study of tls deployment. In *Proceedings of the Internet Measurement Conference 2018*, pages 415–428. ACM, 2018.
- [36] Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. Neural architectures for named entity recognition. *arXiv preprint arXiv:1603.01360*, 2016.

- [37] Martha K Landesberg, Toby Milgrom Levin, Caroline G Curtin, and Ori Lev. Privacy online: A report to congress. NASA, (19990008264), 1998.
- [38] A. Le, J. Varmarken, S. Langhoff, A. Shuba, M. Gjoka, and A. Markopoulou. AntMonitor: A System for Monitoring from Mobile Devices. In *Proc. of ACM SIGCOMM C2B(1)D Workshop*, 2015.
- [39] Jaeho Lee and Dan S Wallach. Removing secrets from android’s tls. In *Proc. Network and Distributed System Security Symposium (NDSS)*, 2018.
- [40] Ilias Leontiadis, Christos Efstratiou, Marco Picone, and Cecilia Mascolo. Don’t kill my ads!: balancing privacy in an ad-supported mobile application market. In *Proceedings of the Twelfth Workshop on Mobile Computing Systems & Applications*, page 2. ACM, 2012.
- [41] Yibin Li, Wenyun Dai, Zhong Ming, and Meikang Qiu. Privacy protection for preventing data over-collection in smart city. *IEEE Transactions on Computers*, 65(5):1339–1350, 2015.
- [42] Jialiu Lin, Shahriyar Amini, Jason I Hong, Norman Sadeh, Janne Lindqvist, and Joy Zhang. Expectation and purpose: understanding users’ mental models of mobile app privacy through crowdsourcing. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, pages 501–510. ACM, 2012.
- [43] Martina Lindorfer, Matthias Neugschwandtner, Lukas Weichselbaum, Yanick Fratantonio, Victor Van Der Veen, and Christian Platzer. Andrubis–1,000,000 apps later: A view on current android malware behaviors. In *2014 Third International Workshop on Building Analysis Datasets and Gathering Experience Returns for Security (BADGERS)*, pages 3–17. IEEE, 2014.
- [44] Minxing Liu, Haoyu Wang, Yao Guo, and Jason Hong. Identifying and analyzing the privacy of apps for kids. In *Proceedings of the 17th International Workshop on Mobile Computing Systems and Applications*, pages 105–110. ACM, 2016.
- [45] Xiaohua Liu, Shaodian Zhang, Furu Wei, and Ming Zhou. Recognizing named entities in tweets. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pages 359–367. Association for Computational Linguistics, 2011.
- [46] Mary Madden, Amanda Lenhart, Sandra Cortesi, Urs Gasser, Maeve Duggan, Aaron Smith, and Meredith Beaton. Teens, social media, and privacy. *Pew Research Center*, 21:2–86, 2013.
- [47] Davide Maiorca, Davide Ariu, Iginio Corona, Marco Aresu, and Giorgio Giacinto. Stealth attacks: An extended insight into the obfuscation effects on android malware. *Computers & Security*, 51:16–31, 2015.
- [48] Kirsten Martin. Transaction costs, privacy, and trust: The laudable goals and ultimate failure of notice and choice to respect privacy online. 2013.
- [49] Aleecia M McDonald and Lorrie Faith Cranor. The cost of reading privacy policies. *ISJLP*, 4:543, 2008.
- [50] George R Milne and Mary J Culnan. Strategies for reducing online privacy risks: Why consumers read (or don’t read) online privacy notices. *Journal of interactive marketing*, 18(3):15–29, 2004.
- [51] George R Milne, Mary J Culnan, and Henry Greene. A longitudinal assessment of online privacy notice readability. *Journal of Public Policy & Marketing*, 25(2):238–249, 2006.
- [52] Jonathan A Obar and Anne Oeldorf-Hirsch. The biggest lie on the internet: Ignoring the privacy policies and terms of service policies of social networking services. *Information, Communication & Society*, pages 1–20, 2018.
- [53] Elleen Pan, Jingjing Ren, Martina Lindorfer, Christo Wilson, and David Choffnes. Panoptispy: Characterizing audio and video exfiltration from android applications. *Proceedings on Privacy Enhancing Technologies*, 2018(4):33–50, 2018.
- [54] Harshvardhan J Pandit, Declan O’Sullivan, and Dave Lewis. Queryable provenance metadata for gdpr compliance, 2018.
- [55] Irene Pollach. What’s wrong with online privacy policies? *Communications of the ACM*, 50(9):103–108, 2007.
- [56] Ashwin Rao, Justine Sherry, Arnaud Legout, Arvind Krishnamurthy, Walid Dabbous, and David Choffnes. Meddle: middleboxes for increased transparency and control of mobile traffic. In *CoNEXT Student Workshop*, 2012.
- [57] A. Razaghpanah, N. Vallina-Rodriguez, S. Sundaresan, C. Kreibich, P. Gill, M. Allman, and V. Paxson. Haystack: In Situ Mobile Traffic Analysis in User Space. *arXiv preprint arXiv:1510.01419*, 2015.
- [58] Abbas Razaghpanah, Arian Akhavan Niaki, Narseo Vallina-Rodriguez, Srikanth Sundaresan, Johanna Amann, and Phillipa Gill. Studying tls usage in android apps. In *Proceedings of the 13th International Conference on emerging Networking EXperiments and Technologies*, pages 350–362. ACM, 2017.

- [59] Abbas Razaghpanah, Rishab Nithyanand, Narseo Vallina-Rodriguez, Srikanth Sundaresan, Mark Allman, Christian Kreibich, and Phillipa Gill. Apps, trackers, privacy, and regulators: A global study of the mobile tracking ecosystem. 2018.
- [60] Abbas Razaghpanah, Narseo Vallina-Rodriguez, Srikanth Sundaresan, Christian Kreibich, Phillipa Gill, Mark Allman, and Vern Paxson. Haystack: In situ mobile traffic analysis in user space. *ArXiv e-prints*, 2015.
- [61] Irwin Reyes, Primal Wijesekera, Joel Reardon, Amit Elazari Bar On, Abbas Razaghpanah, Narseo Vallina-Rodriguez, and Serge Egelman. "Won't Somebody Think of the Children?" Examining COPPA Compliance at Scale. *Proceedings on Privacy Enhancing Technologies*, 2018(3):63–83, 2018.
- [62] Alan Ritter, Sam Clark, Oren Etzioni, et al. Named entity recognition in tweets: an experimental study. In *Proceedings of the conference on empirical methods in natural language processing*, pages 1524–1534. Association for Computational Linguistics, 2011.
- [63] Xinguang Sheng and Lorrie Faith Cranor. An evaluation of the effect of us financial privacy legislation through the analysis of privacy policies. *ISJLP*, 2:943, 2005.
- [64] Anastasia Shuba, Anh Le, Minas Gjoka, Janus Varmarken, Simon Langhoff, and Athina Markopoulou. Antmonitor: Network traffic monitoring and real-time prevention of privacy leaks in mobile devices. In *Proceedings of the 2015 Workshop on Wireless of the Students, by the Students, & for the Students*, pages 25–27. ACM, 2015.
- [65] Rocky Slavin, Xiaoyin Wang, Mitra Bokaei Hosseini, James Hester, Ram Krishnan, Jaspreet Bhatia, Travis D Breaux, and Jianwei Niu. Toward a framework for detecting privacy policy violations in android application code. In *Proceedings of the 38th International Conference on Software Engineering*, pages 25–36. ACM, 2016.
- [66] Robert H Sloan and Richard Warner. Beyond notice and choice: Privacy, norms, and consent. *J. High Tech. L.*, 14:370, 2014.
- [67] FTC Staff. Protecting consumer privacy in an era of rapid change—a proposed framework for businesses and policymakers. *Journal of Privacy and Confidentiality*, 3(1), 2011.
- [68] U.S. Federal Trade Commission. FTC Warns Children's App Maker BabyBus About Potential COPPA Violations, 2014.
- [69] U.S. Federal Trade Commission. Complying with COPPA: Frequently Asked Questions, 2015.
- [70] U.S. Federal Trade Commission. Mobile Advertising Network InMobi Settles FTC Charges It Tracked Hundreds of Millions of Consumers' Locations Without Permission, 2016.
- [71] U.S. Federal Trade Commission. Children's online privacy protection rule: A six-step compliance plan for your business. <https://www.ftc.gov/tips-advice/business-center/guidance/childrens-online-privacy-protection-rule-six-step-com> June 2017. Accessed: November 30, 2017.
- [72] Junia Valente and Alvaro A Cardenas. Security & privacy in smart toys. In *Proceedings of the 2017 Workshop on Internet of Things Security and Privacy*, pages 19–24. ACM, 2017.
- [73] Narseo Vallina-Rodriguez, Jay Shah, Alessandro Finamore, Yan Grunenberger, Konstantina Papagiannaki, Hamed Haddadi, and Jon Crowcroft. Breaking for commercials: characterizing mobile advertising. In *Proceedings of the 2012 Internet Measurement Conference*, pages 343–356. ACM, 2012.
- [74] Liesbet Van Zoonen. Privacy concerns in smart cities. *Government Information Quarterly*, 33(3):472–480, 2016.
- [75] Primal Wijesekera, Arjun Baokar, Ashkan Hosseini, Serge Egelman, David Wagner, and Konstantin Beznosov. Android permissions remystified: A field study on contextual integrity. In *24th {USENIX} Security Symposium ({USENIX} Security 15)*, pages 499–514, 2015.
- [76] Shomir Wilson, Florian Schaub, Aswarth Abhilash Dara, Frederick Liu, Sushain Cherivirala, Pedro Giovanni Leon, Mads Schaarup Andersen, Sebastian Zimmeck, Kanthashree Mysore Sathyendra, N Cameron Russell, et al. The creation and analysis of a website privacy policy corpus. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1330–1340, 2016.
- [77] Shomir Wilson, Florian Schaub, Rohan Ramanath, Norman Sadeh, Fei Liu, Noah A Smith, and Frederick Liu. Crowdsourcing annotations for websites' privacy policies: Can it really work? In *Proceedings of the 25th International Conference on World Wide Web*, pages 133–143. International World Wide Web Conferences Steering Committee, 2016.

- [78] Benjamin Yankson, Farkhund Iqbal, and Patrick CK Hung. Privacy preservation framework for smart connected toys. In *Computing in Smart Toys*, pages 149–164. Springer, 2017.
- [79] Le Yu, Tao Zhang, Xiapu Luo, and Lei Xue. Autoppg: Towards automatic generation of privacy policy for android applications. In *Proceedings of the 5th Annual ACM CCS Workshop on Security and Privacy in Smartphones and Mobile Devices*, pages 39–50. ACM, 2015.
- [80] Sebastian Zimmeck, Ziqi Wang, Lieyong Zou, Roger Iyengar, Bin Liu, Florian Schaub, Shomir Wilson, Norman Sadeh, Steven Bellovin, and Joel Reidenberg. Automated analysis of privacy requirements for mobile apps. In *2016 AAAI Fall Symposium Series*, 2016.
- [81] Sebastian Zimmeck, Ziqi Wang, Lieyong Zou, Roger Iyengar, Bin Liu, Florian Schaub, Shomir Wilson, Norman Sadeh, Steven M Bellovin, and Joel Reidenberg. Automated analysis of privacy requirements for mobile apps. In *24th Network & Distributed System Security Symposium (NDSS 2017)*, NDSS, 2017.