"Is Our Children's Apps Learning?" Automatically Detecting COPPA Violations

Primal Wijesekera[1,2], Abbas Razaghpanah[3], Joel Reardon[1,4], Irwin Reyes[4]
Narseo Vallina-Rodriguez[4,5], Serge Egelman[1,4], Christian Kreibich[4,6]
[1]UC Berkeley, [2]UBC, [3]Stony Brook University, [4]ICSI, [5]IMDEA Networks, [6] Lastline

ABSTRACT

In recent years, a market of games and learning apps for children has flourished in the mobile world. Many of these often "free" mobile apps have access to a variety of sensitive personal information about the user, which the app author can leverage to increase revenue via advertising or other means. In the United States, the Children's Online Privacy Protection Act (COPPA) protects children's privacy, requiring parental consent to the use of personal information and prohibiting behavioral advertising and online tracking.

In this work, we present our ongoing effort to develop a method to automatically evaluate mobile apps' COPPA compliance. Our method combines dynamic execution analysis (to track sensitive resource access at runtime) with traffic monitoring (to reveal private information leaving the device and recording with whom it gets shared, even if encrypted). We complement our empirical technical observations with legal analysis of the apps' corresponding privacy policies.

As a proof of concept, we scrape the Google Play store for apps that declare their target group to be less than 13 years of age, which subjects them to COPPA's regulations. We automate app execution on an instrumented version of the Android OS, recording the apps' access to and transmission of sensitive information. To contextualize third parties (e.g., advertising networks) with whom the apps share information, we leverage a crowdsourced dataset collected by Haystack, our Android-based device-local traffic inspection platform. Our effort illuminates apps' compliance with COPPA and catalogs the organizations that collect sensitive user information. We find several likely COPPA violations in our preliminary corpus, including omission of prior consent and active sharing of persistent identifiers with third-party services for tracking and profiling of children.

1. INTRODUCTION

Recent years have seen significant increase in smartphone use among children [27]. Accordingly, a large number of mobile games and educational applications ("apps") have been developed for use by children thanks to the ubiquitous nature of mobile platforms and the usability improvements introduced by modern touch screens. While users can typically download these apps free of cost, they often generate revenue through advertising [32]—including some business models that tailors ads to users' interests by tracking their online behavior over time, including access to users' personal data stored on their devices (e.g., contact lists, location trails, or the web browsing history).

Previous work has documented apps using personal information in ways unexpected or not apparent to their users [19]. While such privacy violations prove worrisome for anyone, children are particularly vulnerable due to their inability to understand the importance of personal information and to provide informed consent. The United States recognizes the lasting effect that privacy violations may have on children, and has passed strong legislation—the Children's Online Privacy Protection Act (COPPA), enforced by the Federal Trade Commission (FTC)—to regulate how web sites and mobile apps can collect, and share with third parties, private information of children under the age of 13 [20]. COPPA rules require verified parental consent prior to collection of Personally Identifiable Information (PII), and that services take steps to ensure that the consenting party is in fact a legal parent or guardian.

Despite regulatory efforts to protect sensitive audiences, the current status of mobile apps' compliance with COPPA rules remains largely unknown. Prior research by FTC staff involved laboriously downloading popular children's apps and manually examining them. In one report, the researchers uncovered numerous violations [29]. In a follow-up study performed almost a year later, they found little progress with regard to COPPA compliance [30]. Since both studies involved manual evaluation of apps, they covered only a small subset of available children's apps and looked for only a subset of possible COPPA violations. It also remains unclear whether anything has changed in the intervening four years, due to the continued threat of sanctions for violators.

In this paper we present our ongoing effort to build a method for analyzing apps' COPPA compliance at scale. Our goal is to increase transparency by draw-

ing attention to apps' sensitive data usage and sharing practices, especially as it concerns the data of children. Our method combines dynamic analysis of Android app behaviors during runtime [33] with in-depth inspection of network traffic [28] to analyze how apps access and share sensitive personal information. Our method records whether an app engages in tracking activity, whether it discloses this to the end user, whether it shares personal data with third parties, and whether it asks for parental consent, as required by the COPPA statute. We complement our empirical analysis on the technical side with a method to extract and analyze if the privacy policies available on Google Play inform users of potential tracking activities. Our preliminary results reveal several potential COPPA violations, including apps accessing PII without prior consent and actively sharing persistent identifiers with third-party services that enable the tracking and profiling of children across different Internet services.

2. LEGAL PROTECTIONS

In 1998, the United States Congress first enacted the Children's Online Privacy Protection Act (COPPA) and amended it in 2012 to add new categories to the definition of PII. It aims to protect children under the age of 13 who use commercial websites, online games, and mobile apps [20]. The main objective of COPPA is to give parents control over how vendors access their children's personal information and the organizations receiving such sensitive information.

COPPA has two requirements to help parents make decisions about their children's data when installing a new app: (i) vendors must disclose their PII collection practices (i.e., what are the types of data they access and with whom do they share this data), (ii) vendors must ask for *verifiable parental consent* before first accessing any PII. Information considered PII by the COPPA rules [31] includes first and last name, physical addresses, user account names, phone numbers, social security numbers, device identifiers (such as IMEI, IMSI, MAC addresses and serial numbers), media (such as photos, video, or audio recordings) featuring the child, and precise geolocation information.

COPPA prohibits any form of online tracking for children under the age of 13, including sharing with third-party services such as ad networks and analytics services. The FTC enforces COPPA rules and over the past few years has brought several successful actions against COPPA violators for reasons including not seeking parental consent before accessing PII and sharing persistent identifiers with third-party services [9, 10, 12, 16, 17, 23]. The FTC has so far scrutinized select apps based on complaints or other suspicious behavior reported by the public. Our work intends to understand the extent of compliance among all apps—not just ones reported by the general public—using an automated detection process. We hope that this tool will nudge app developers towards greater compliance.

While COPPA jurisdiction only applies to apps marketed to users in the United States, other countries have their own laws and guidelines to protect children. Canada has different regulations at the provincial and federal level. Federally, it prohibits tracking children across Internet services [13]. Some provincial regulations go further, banning all advertising to children under 13 [5]. The EU is currently adopting a new law regulating children's privacy across all member countries. The new law, *Article 8* of the European Convention on Human Rights [8], mainly focuses on forcing apps to seek parental consent before accessing any PII from children.

3. INDUSTRY RESPONSE

COPPA excludes platforms, hosting services, and distribution channels from any liability: the final product vendor (i.e., the app developer) bears responsibility for complying with the law. Nevertheless, both the Google Play Store and Apple App Store have introduced measures to force app developers to comply with the law; non-compliant apps risk de-listing from the stores.

The Google Play Store introduced specific age categories under the "Designed for Families" program [6], aiming to help parents filter out inappropriate apps. App developers wishing to participate in this program—listing their apps under the Play Store's "Families" category and its under-13 age subcategories—must comply with Google's guidelines for age-appropriate content and advertising, including COPPA compliance. Participating apps must have an ESRB, a content rating [7], rating of "Everyone" (or equivalent), ensure that in-app ads remain appropriate for the target audience, and post a privacy policy on the app's store listing. Developers agree to abide by these standards as long as their apps appear in the "Family" category. No automated system appears to be in place, however, to verify continued compliance after the initial acceptance into the "Designed for Families" program [11].

Similarly, the Apple App Store introduced a special "Kids Category" for children apps. Any developer who wants to list their app in this category must also follow extra policies [2] based on COPPA. Apple has also introduced a family sharing disclosure [3], giving parents more control over the types of data that a children's app can access.

This work focuses on children's apps available through the Google Play Store. The Google Play Store does not automatically classify which submitted apps are family-friendly or directed at young children. Instead, app developers and publishers must self-report children's apps during the app publication process. By having their

apps listed in the "Designed for Families" program and the relevant age subcategories, app developers acknowledge that their app targets users under the age of 13 and therefore makes them liable for any COPPA violations.

4. RELATED WORK

Beyond the two studies performed by the FTC to gauge COPPA compliance [29, 30], previous work in this field has focused primarily on privacy violations of the adult population. Previous work has shown that apps' access to sensitive user data often defies expectations [19, 33]. Researchers have also shown the ineffectiveness of the different privacy regulation models deployed in Android [21, 33].

A study conducted by Liu *et al.* [25] identified almost 68,000 children apps from a set of one million Android apps. They presented a method to identify potential COPPA violations using app metadata publicly available from the apps' public profiles. The study provided no insights into app runtime behaviors or the actual privacy leaks caused by either the apps or organizations behind them.

A study conducted by Hu *et al.* [24] predicted the age target of apps by using app metadata, so as to give parents guidance when selecting apps for their children. While the nature of the content is important for kids' apps, the study did not consider how apps comply with privacy regulations. In contrast, our work examines COPPA compliance among apps that are specifically targeted at kids.

5. THE COPPA COMPLIANCE TESTBED

We now describe our COPPA compliance testbed, which automates the technical analysis of Android apps' COPPA compliance. Our testbed has four broad goals: it (i) identifies children's apps that access sensitive information, (ii) reveals the third parties with whom they share such information, (iii) checks whether the apps request parental consent at runtime, and (iv) assists legal analysts in gauging the extent to which such privacy policies prove informative and correct. We use this testbed to evaluate apps submitted under Google's "Designed for Families" program, as well as those designed for general audiences.

Our testbed consists of LG Nexus 5 phones running a customized version of the Android Open Source Project (AOSP) 6.0.1 Marshmallow [4]. Our instrumentation combines dynamic execution tracing and network traffic analysis, as follows. At runtime, our customized kernel records apps' access to sensitive resources controlled by Android's permissions system, including geolocation data, stored pictures, SMS, browsing history, and media capture (i.e., audio, photos, and video) [33]. Our instrumentation tracks all COPPA-relevant resource requests by monitoring sensitive function calls invoked by the apps under investigation. In addition, it records a host of contextual information surrounding each request, such as the visibility (i.e., foreground or background) of the app requesting the resource. This instrumentation operates at the platform level, allowing us to run and analyze apps from the Google Play Store as-is, without any modification or preprocessing.

To complement the OS-level instrumentation we simultaneously run our ICSI Haystack traffic monitor [28], an Android app freely available via Google Play [22] that helps users understand how their apps transmit private information, including the nature of sensitive data transmitted by mobile apps as well as the recipients with whom the apps share the data (e.g., analytics services and ad networks). Haystack leverages Android's VPN permission to capture and analyze network traffic in user space, on the device. Haystack also intercepts and decrypts data transmitted over TLS, via an optional local TLS interception proxy that we enable for the COPPA analysis.

Haystack benefits our testbed in three ways: (i) it determines whether any COPPA-restricted personal data actually gets transmitted to third parties, (ii) the data provided by Haystack's user-base helps us catalog and label the third-party tracker landscape, allowing us to gauge the role of third-party trackers found on children's apps, and (iii) Haystack complements the OS instrumentation by also identifying privacy leaks that do *not* require explicit Android permissions.

5.1 Automated Testing

We conduct automated testing of apps using the Android Application Exerciser Monkey [18]. This allows us to automate the execution of apps with minimal human intervention. The Monkey naively generates a pseudorandom stream of taps, swipes, button presses, and other simulated input events, which we run for approximately ten minutes. This allows us to explore the app's behavior and observe if any sensitive information leaves the device. After each experiment, we record log data from the resource-access instrumentation and Haystack, as well as the random seed used for the Monkey sequence for debugging and replication. During these experiments, we also take screenshots of the first 30 seconds of app execution, likely to contain any consent form as it should be disclosed before engaging on any user tracking activity, before providing any inputs. We use these images later to identify whether—as mandated by COPPA—parental consent and privacy disclosures appear on apps' landing screens.

5.2 Supervised Analysis

Although the Exerciser Monkey generates useful data for initial analysis, unguided exploration does not result in complete coverage of the app's functionality space.

Multi-step UI elements like text entry boxes (e.g., login) and slider widgets impede the Monkey's progress through an app. COPPA-restricted data, such as audio recordings and photos, often are accessed through similarly complex UIs. The Monkey is unlikely to randomly generate the correct sequence of input events to activate such multi-step UI elements and progress through the app within the allocated run time.

In order to address the Monkey's practical shortcomings, we will recruit human testers to explore apps in a more guided and realistic manner. Our testers will be instructed to activate all the interactive UI elements they see while interacting with each app. For apps in our corpus that have such functionality, we will also ask the testers to record audio and take photos and videos. Testers are given personas with names, email addresses, and COPPA-protected personal information to provide to any apps that request these. We will conduct this human-powered testing on the same hardware and software environment as our automated exploration, and subsequently collect, compare and analyze the same log data. This will also allow us to assess the accuracy and coverage that the Monkey provides.

5.3 Privacy Policy Analysis

The act of collecting certain types of private children's data does not necessarily constitute a COPPA violation. Because collection is permitted provided that the privacy policy discloses it and it happens for an allowable purpose, we must scrutinize the privacy policies for each tested app.

Automated analysis of privacy policies using text mining techniques is not suitable due to their complexity. To answer questions about what provisions of COPPA apply to an app, we will recruit law students from our institution to code the policies from our corpus of apps. This coding will allow us to determine which apps disclose that they collect private information, use persistent identifiers, and so on. With this information, we can identify the apps for which we are certain that particular behaviors must not be performed. Moreover, by having multiple law students—as well as laypeople—code the policies, we can compare their results for consistency. Recruiting laypeople will allow us to examine how well a policy expresses various practices to potential users (or rather, their parents).

To simplify this privacy policy coding, we scraped the Google Play Store for the privacy policies of all the apps in our corpus. We built a tool that presents users with a particular privacy policy along with questions about it. Participants may further select relevant parts of the policy to augment their multiple-choice answers with examples. Once policies are coded in this manner, we will be able to automatically determine which observed practices are being disclosed, and which present privacy violations.

Permission	Declared	Used
ACCESS_COARSE_LOCATION	2	0
ACCESS_FINE_LOCATION	1	0
ACCESS_WIFI_STATE	7	3
READ_CALL_LOG	0	0
READ_SMS	0	0
SEND_SMS	0	0

Table 1: Instrumented permissions declared and used by a random selection of 22 corpus apps

6. PRELIMINARY RESULTS

We report on two sets of analyses; one focusing on apps examined in the testbed, the second investigating COPPA violations in the broader datasets reported to us by users of the Haystack app.

6.1 Testbed-driven Analysis

We began our experiments by downloading mobile apps that are targeted for ages 13 and under in Google Play's Family category. We reiterate that for an app to be listed under this category, the app developer acknowledged that the app is in fact suitable for that age group. COPPA refers to this as *actual knowledge* on behalf of the developer that children under the age of 13 will use this app, rendering the developer liable for possible violations.

As proof of concept we performed a small-scale analysis of 25 apps drawn randomly from a corpus of 620 apps. Of these, three did not declare the INTERNET permission, which is required for apps to open network sockets. Because it is unlikely that these apps communicate with remote servers through other channels, we excluded them from further analysis. Table 1 summarizes the number of apps that declare permissions that our instrumentation analyses, as well as the number of apps observed using these permissions during a ten-minute Monkey run.

For the most part, the apps we evaluated in this small-scale test did not access sensitive resources monitored by our instrumented platform. The most frequently declared permission, ACCESS_WIFI_STATE, is primarily used to check for a Wi-Fi Internet connection. It could also serve to retrieve saved routers' BSSIDs—a proxy for location information [26]. Two apps declared ACCESS_COARSE_LOCATION, which allows for the scanning of all Wi-Fi-router BSSIDs in range (rather than just the BSSID of the connected network). Of these apps, one further declared ACCESS_FINE_LOCATION, for full geolocation using GPS and cellular towers. Neither of these apps were observed using these declared location permissions (i.e., the methods to retrieve the actual location data were never called). As mentioned

Action	Adventure	Arcade	Board
Casual	Education	Educational	Personalization
Puzzle	Racing	Role Playing	Simulation
Strategy			

Table 2: Selected app categories that can potentially be used by children.

earlier, the Monkey does not provide complete coverage of the app space, so it is possible that these functions simply were not triggered by the random input stream. Further testing with supervised exploration of apps is warranted.

We also checked the use and sharing of persistent identifiers with third parties. Google recommends its Android Advertising ID (AAID) as the *only* persistent device-level identifier for tracking and marketing purposes [1]. The AAID is accessible to any installed app, without special permissions, and remains constant unless the user manually regenerates it through the settings UI or restores the phone to factory settings. Using Haystack to examine the contents of a subset of HTTPS GET and POST requests, we identified two apps sharing our testbed's AAID to multiple third parties: one shared it with three analytics platforms and the other with four advertising platforms. In all these cases, the requests also included the app's identifying package name, which can be used to associate a persistent ID with the use of child-targeted apps (i.e., the recipients would know that the identifiers originated from apps targeted at children).

6.2 Haystack Dataset

Next, we mined anonymized traffic traces collected from 690 Haystack users for potential COPPA violations.[1] This dataset complements the artificial UI events generated by our testbed with traffic monitored in-situ on real-world user activity.

To widen focus from the explicitly child-targeting apps in the "Designed for Families" program, we stipulate that children will nevertheless also often explore games and similarly interesting apps with no maturity rating (i.e., an ESRB rating of "Everyone"), and specifically focus on such apps. We focus our analysis on unique identifiers (e.g., IMEI, IMSI, MAC addresses, and serial numbers) leaked by mobile apps with no maturity constraint and belonging to the categories listed in Table 2.

Our analysis revealed 18 games and two personalization apps that cause potential COPPA violations by sharing unique identifiers with 15 third-party services. In order to collect the device MAC address and serial number—two unique identifiers with the same privacy impact as the IMEI and IMSI values—app developers

[1] We refer the reader to prior work on Haystack for details on the platform, our data anonymization process, and IRB considerations [28].

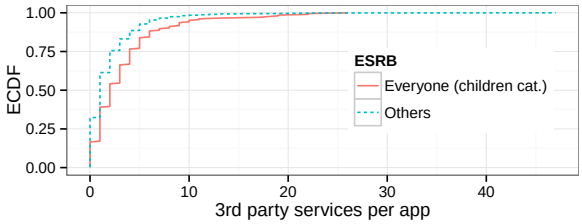

Figure 1: CDF of the number of trackers for general audience apps in app categories defined in Table 2 ($N = 205$) and the apps falling in other categories ($N = 1646$).

do not need to request any specific permission [28]: this information is accessible by invoking an undocumented system-maintained command (`getprop`), which contains different device properties and system configuration values. This suggests that app developers are deliberately attempting to track users without their awareness and consent. Finally, ten of these apps upload sensitive unique identifiers over unencrypted channels, thus easing user profiling by observers of network traffic. (Transmitting PII over unencrypted channels is itself a potential COPPA violation.)

We conclude our analysis with a comparison of the number of trackers found in apps in the categories listed in Table 2 with the number of trackers in the apps falling into any other category. To this end, we leverage the list of domains associated with third-party services produced by the ICSI Haystack team [15]. As we can see in Figure 1, despite the difference in the number of apps in each category, apps that may be used by children tend to have a higher number of trackers than other apps. Around 80% of the apps potentially used by children use at least one tracking service, as opposed to 65% of the apps falling in other app categories. Our analysis identified 7 games reaching more than 15 third-party trackers. After inspecting their Google profiles manually, we observed that these are popular children games (not listed in the Family categories) with more than 100 million installs and with positive ratings (4+ stars) implemented by game developers awarded with the "Top Developer" badge in Google Play [14].

7. CONCLUSIONS

This paper presents a first look at our COPPA compliance testbed, which uniquely combines dynamic execution tracing of Android apps, real-time network traffic analysis, and human-analyst feedback on applicable privacy policies to produce app-specific profiles of potential COPPA violations in apps targeting children.

Our preliminary analysis of apps on the Google Play Store finds strong evidence of apps explicitly targeted at children sending private information to third-party services and advertisers.

8. REFERENCES

[1] Advertising ID - Developer Console Help. https://support.google.com/googleplay/android-developer/answer/6048248?hl=en.

[2] App Store Review Guidelines. https://developer.apple.com/app-store/review/guidelines/.

[3] Apple ID and Family Sharing Disclosure. http://www.apple.com/legal/privacy/en-ww/parent-disclosure/.

[4] Codenames, Tags, and Build Numbers — Android Open Source Project. https://source.android.com/source/build-numbers.html.

[5] Consumer Protection Act. http://legisquebec.gouv.qc.ca/en/showdoc/cs/P-40.1.

[6] Developer Policy Center – Families. https://play.google.com/about/families/designed-for-families/.

[7] Entertainment Software Rating Board. http://www.esrb.org/ratings/ratings_guide.aspx.

[8] European Convention on Human Rights - Article 8. http://echr-online.info/article-8-echr/.

[9] FTC announces first mobile app case. https://www.ftc.gov/news-events/blogs/business-blog/2011/08/ftc-announces-first-mobile-app-case.

[10] FTC Settles with Children's Gaming Company For Falsely Claiming To Comply With International Safe Harbor Privacy Framework.

[11] Google Play for Families FAQ — Android Developers. https://developer.android.com/distribute/googleplay/families/faq.html.

[12] Mobile Advertising Network InMobi Settles FTC Charges. https://www.ftc.gov/news-events/press-releases/2016/06/mobile-advertising-network//-inmobi-settles-ftc-charges-it-tracked.

[13] Privacy and kids. https://www.priv.gc.ca/en/privacy-topics/privacy-and-kids/.

[14] The Google Play Opportunity. https://developer.android.com/distribute/googleplay/about.html.

[15] The ICSI Haystack panopticon. https://www.haystack.mobi/panopticon.

[16] Two App Developers Settle FTC Charges For Sharing Persistent Identifiers. https://www.ftc.gov/news-events/press-releases/2015/12/two-app-developers-settle-ftc-charges-they-violated-childrens.

[17] Yelp, TinyCo Settle FTC Charges Their Apps Improperly Collected Children's Personal Information. https://www.ftc.gov/news-events/press-releases/2014/09/yelp-tinyco-settle-ftc-charges-their-apps-improperly-collected.

[18] Android Developer's Documentation. Android developers: UI/application exerciser monkey. http://developer.android.com/tools/help/monkey.html.

[19] W. Enck, P. Gilbert, B.-G. Chun, L. P. Cox, J. Jung, P. McDaniel, and A. N. Sheth. Taintdroid: an information-flow tracking system for realtime privacy monitoring on smartphones. In *Proceedings of the 9th USENIX Conference on Operating Systems Design and Implementation*, OSDI'10, pages 1–6, Berkeley, CA, USA, 2010. USENIX Association.

[20] Federal Trade Comission. Children's Online Privacy Protection Rule ("COPPA"), 1998. https://www.ftc.gov/enforcement/rules/rulemaking-regulatory-reform-proceedings/childrens-online-privacy-protection-rule.

[21] A. P. Felt, E. Ha, S. Egelman, A. Haney, E. Chin, and D. Wagner. Android permissions: user attention, comprehension, and behavior. In *Proc. of the Eighth Symposium on Usable Privacy and Security*, SOUPS '12, New York, NY, USA, 2012. ACM.

[22] Google Play. Icsi haystack. https://play.google.com/store/apps/details?id=edu.berkeley.icsi.haystack&hl=en.

[23] S. Gressin. COPPA: When Persistence Doesn't Pay. https://www.ftc.gov/news-events/blogs/business-blog/2015/12/coppa-when-persistence-doesnt-pay, December 17 2015.

[24] B. Hu, B. Liu, N. Z. Gong, D. Kong, and H. Jin. Protecting your children from inappropriate content in mobile apps: An automatic maturity rating framework. In *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management*, pages 1111–1120. ACM, 2015.

[25] M. Liu, H. Wang, Y. Guo, and J. Hong. Identifying and analyzing the privacy of apps for kids. In *ACM HotMobile*, 2016.

[26] C. Matte and M. Cunche. DEMO: Panoptiphone: How Unique is Your Wi-Fi Device? In *ACM WiSec 2016*.

[27] C. S. Media. Zero to Eight: Children's Media Use in America 2013. https://www.commonsensemedia.org/sites/default/files/research/zero-to-eight-2013.pdf, 2013.

[28] 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 e-prints*, 2015.

[29] U.S. Federal Trade Commission. Mobile Apps for Kids: Current Privacy Disclosures are Disappointing. http://www.ftc.gov/os/2012/02/120216mobile_apps_kids.pdf, February 2012.

[30] U.S. Federal Trade Commission. Mobile Apps for Kids: Disclosures Still Not Making the Grade. https://www.ftc.gov/reports/mobile-apps-kids-disclosures-still-not-making-grade, December 2012.

[31] 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-compliance, June 2013.

[32] N. Vallina-Rodriguez, J. Shah, A. Finamore, Y. Grunenberger, K. Papagiannaki, H. Haddadi, and J. Crowcroft. Breaking for commercials: characterizing mobile advertising. In *ACM IMC*, 2012.

[33] P. Wijesekera, A. Baokar, A. Hosseini, S. Egelman, D. Wagner, and K. Beznosov. Android permissions remystified: A field study on contextual integrity. In *24th USENIX Security Symposium (USENIX Security 15)*, pages 499–514, Washington, D.C., Aug. 2015. USENIX Association.