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

Irwin Reyes[4], Primal Wijesekera[1,2], Abbas Razaghpanah[3], Joel Reardon[1,4], 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 app developers can monetize 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 empirical technical observations with legal analysis of the apps' corresponding privacy policies.

As a proof of concept, we scraped the Google Play store for apps distributed in categories specifically targeting users under than 13 years of age, which subjects these products to COPPA's regulations. We automated 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 leveraged a crowdsourced dataset collected by the Lumen Privacy Tool (formerly Haystack) [27], an Android-based device-local traffic inspection platform. Our effort seeks to illuminate apps' compliance with COPPA and catalog the organizations that collect sensitive user information. In our preliminary results, we find several likely COPPA violations, including omission of prior consent and active sharing of persistent identifiers with third-party services for tracking and profiling of children. These results demonstrate our testbed's capability to detect different types of possible violations in the market for children's apps.

I. Introduction

While the European Union has a data protection directive [19] that protects consumers broadly across many industries, no such legislation exists in the United States. Instead, there are sector-specific laws that govern how personal data may be used by businesses in that sector. One sector that has stringent requirements and that sees frequent enforcement is the sector that involves children. 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 private information of children under the age of 13 [20] as well as when such collected data may be shared with third parties. COPPA rules require *verified parental consent* prior to collection of any Personally Identifiable Information (PII), and that services take steps to ensure that the consenting party is in fact a legal parent or guardian. It is for this reason that many websites targeted at children collect as little PII as possible, so as to not violate COPPA. Until recently, however, enforcement efforts on mobile platforms have been much rarer, partially because the investigation of potential violations has been a laborious process.

Recent years have seen significant increase in smartphone use among children [26]. Accordingly, a large number of mobile games and educational applications ("apps") have been developed for use by children due to the ubiquitous nature of mobile platforms and the usability improvements introduced by modern touch screens. While these apps are often free of cost, they may generate revenue through advertising [31]—including some business models that tailors ads to users' interests by tracking their online behavior over time or accessing personal data stored on the users' devices (e.g., contacts, location trails, or browsing history).

Previous work has documented apps using personal information in ways unexpected or not apparent to their users [18], [33]. 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.

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 [28]. In a follow-up study performed almost a year later, they found little improvement with regard to COPPA compliance [29]. 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, despite the continued threat of sanctions for violators.

In this work. we present our ongoing effort to build a method for analyzing apps' COPPA compliance at scale. Our

goal is to increase transparency by drawing 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 [27] to analyze how apps access and share sensitive personal information. Our method records whether an app engages in tracking activity, whether it discloses this tracking to the end user, whether it shares personal data directly to third parties, and whether it asks for parental consent. 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. The use of our automated analysis tools has already brought about positive change: at least one developer removed an advertising library after we brought its behaviors to their attention.

II. 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. COPPA 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) developers 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) developers must ask for *verifiable parental consent* before first accessing any PII. Information considered PII by COPPA [30] 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], [15], [16], [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 our tool pushes app developers towards greater compliance, while also allowing stakeholders—both regulators and parents—to more easily detect violations.

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 regulations at both the provincial and federal level. Federally, it prohibits tracking children across Internet services [13]. Some provinces further ban 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.

III. 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 compliance. Nevertheless, both the Google Play Store and Apple App Store have 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 rating (a content rating for age appropriateness [7]) 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's 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.

We focus 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.

IV. Related Work

Beyond the two studies performed by the FTC to gauge COPPA compliance [28], [29], 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 [18], [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.

V. The COPPA Compliance Testbed

We now describe our testbed, which automates the technical analysis of Android apps for COPPA compliance. Our testbed has four broad goals: (i) to identify children's apps that access sensitive information, (ii) to reveal any third parties with whom they share such information, (iii) to check whether the apps request parental consent at runtime, and (iv) to assist 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, text messages, 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*, i.e., without any modification or preprocessing.

To complement the OS-level instrumentation, we simultaneously run Lumen Privacy Monitor [27]. Lumen, which is freely available via Google Play [22], helps users understand how their apps transmit private information, including the nature of sensitive data transmitted by mobile apps, as well as the recipients of data shared by the apps (e.g., analytics services and ad networks). Lumen leverages Android's VPN permission to capture and analyze network traffic in user space, on the device. Lumen also intercepts and decrypts data transmitted over TLS, via an optional local TLS interception proxy that we enable for the COPPA analysis. [1]

Lumen benefits our testbed in three ways: (i) it determines whether any COPPA-restricted personal data actually gets transmitted to third parties, (ii) the crowd-sourced anonymized data provided by Lumen'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) Lumen complements the OS instrumentation by also identifying privacy leaks that do *not* require explicit Android permissions. These include OS build fingerprints, software identifiers like the Android Advertising ID, and hardware identifiers like the Wi-Fi MAC address.

A. Automated Testing

We conduct automated testing of apps using the Android Application Exerciser Monkey [17]. 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 Lumen, as well as the random seed used for the Monkey sequence for debugging and replication.

B. 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 hired human testers from the UC Berkeley undergraduate population to explore apps in a more guided and realistic manner. We instructed our testers to activate all the interactive UI elements they see while interacting with each app. We also asked testers to record audio and take photos and videos for apps that have this functionality. Testers were given personas with names, email addresses, and COPPA-protected personal information to provide to any apps that request them. Human-powered testing is performed on the same hardware and software environment as our automated exploration, and subsequently collects, compares, and analyzes the same log data. Scripts automate the environment's setup and teardown, as well as enforce a 10-minute exploration time limit for each app—the same exploration length as in the automated runs.

[1] We refer the reader to prior work on Lumen for details on the platform, our data anonymization process, and IRB considerations [27].

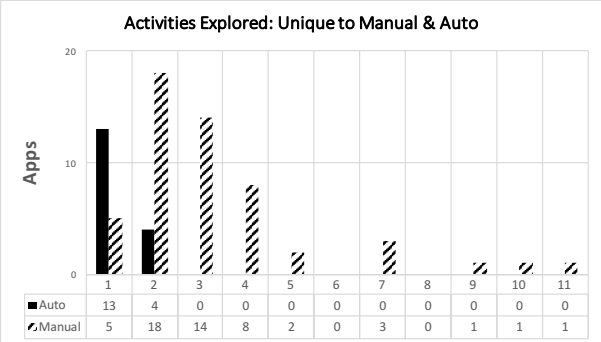

Fig. 1. Apps where Activities were explored only by human testers versus only by the automated system

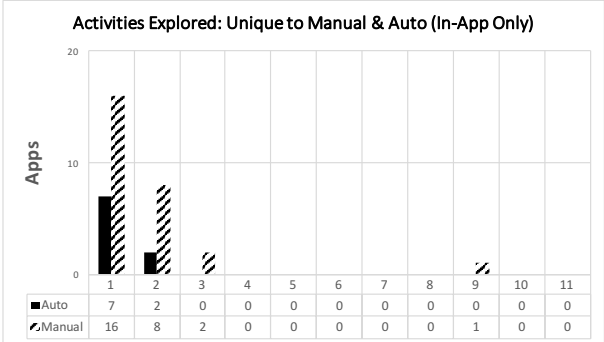

Fig. 2. Apps where Activities were explored only by human testers versus only by the automated system, limited to just Activities that belong to the app being tested

This allows us to assess the accuracy and coverage that the Monkey provides compared to a human user.

VI. Preliminary Results

We report on two sets of analyses; apps we examined in our testbed and a broader dataset reported to us by users of Lumen.

A. Testbed-driven Analysis

As a proof of concept, we conducted a small-scale analysis of 111 apps drawn randomly from our corpus of 446 children's apps. We generated the corpus by scraping app packages from the "Ages 5 & under," "Ages 6–8," and "Ages 9 & up" subcategories in the Google Play Store. We reiterate that for an app to be listed under these categories, 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 the product, rendering the developer liable for possible violations.

1) Automatic and Manual Exploration: Of the 111 apps we evaluated using our automated system, 61 have also been tested under manual human supervision, at the time of this writing. Human-supervised testing establishes a baseline with which to compare the thoroughness of automated testing. Our instrumentation records which Android Activities (i.e., discrete screens and tasks) were visited during the runs. By comparing the names of observed Activities, we can determine which screens were seen only by human testers, and which were seen only by the automated system.

In 53 out of the 61 apps, human-driven testing explored at least one screen not seen in automatic testing (Figure 1). In 17 of the 61 apps, the automatic Exerciser Monkey saw at least one screen human testers did not see. We examined the recorded Activity names to determine why human testers tended to find more screens that the automatic testing missed. In many cases, the ones only seen by human testers actually belonged to external applications such as web browsers, YouTube channels, and app stores. In-app links and buttons can spawn these external Activities. We consider these external Activities to be outside the immediate scope of testing goals: app developers are not responsible for how other companies' programs behave, and those external applications may not have necessarily been distributed through channels explicitly meant for users under 13 years of age.

When discounting Activities not part of the app under test, human exploration yielded screens unvisited by the Exerciser Monkey in 27 apps (Figure 2). The Exerciser Monkey saw screens the human missed in 9 apps. Though both the Exerciser Monkey and the human had comparable coverage with one another in a majority of apps, both missed screens that the other explored for a significant number of cases—more in automated testing than human-directed. This suggests that the two testing modes should complement one another in some cases: automated testing can provide an initial set of observations at scale, then manual testing for a selection of those apps.

2) Android Advertising ID: COPPA restricts the use and collection of persistent identifiers "that can be used to recognize a user over time and across different websites or online services." A variety of persistent identifiers are available on Android devices, such as the Wi-Fi adapter's MAC address, phone IMEI, and SIM card serial numbers.

The Android Advertising ID (AAID) is a persistent identifier generated by the operating system and is visible to all apps. Apps are not required to inform the user or obtain consent before accessing it. The user may, however, opt out of AAID tracking or generate a new AAID using the Android system settings. Google Play best practices recommend that the AAID be used *exclusively* for advertising and analytics purposes [1]. Google also strongly discourages associating the AAID with other device identifiers.

We found that 50 of the 111 apps transmitted the AAID to third parties. These third-party communications were frequently accompanied by other device identifiers such as the IMEI and device ID. Third-party domains that received the AAID along with another identifier include *appsflyer.com*, *tdcv3.talkingdata.net*, and *data.flurry.com*.

We also performed limited testing of how apps and libraries observe user preferences to opt-out of AAID tracking. In one children's app, BabyFirst's "Peekaboo Goes Camping Game," the AAID (along with the device ID) continues to be sent to the *appsflyer.com* third-party domain even after the user opts

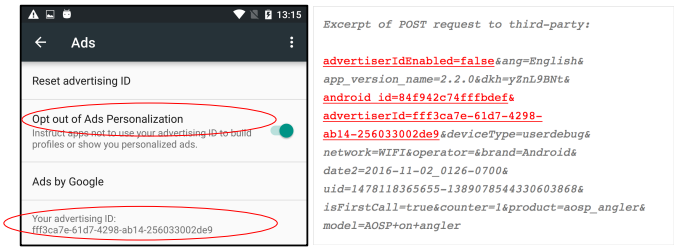

Fig. 3. An app continues to use the AAID even after the user opts out of tracking

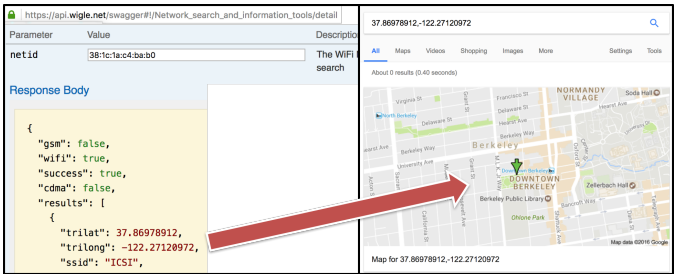

Fig. 4. Locating the ICSI offices using the transmitted MAC address

out of this tracking in the system settings (Figure 3).

We disclosed this finding to BabyFirst weeks before submitting this paper, but we did not receive a response.

3) MAC Geolocation Leakage: Street-resolvable geolocation is prohibited under COPPA. Although app developers and third-party firms can (and often do) use IP address geolocation to identify their users' cities, this data is insufficiently personally-identifiable to run afoul of the regulations. On the other hand, children's apps that share GPS satellite, cell network, and Wi-Fi network geolocation are likely violating the law, as those are accurate within tens of meters and therefore sufficient to determine a residential address with high confidence. To mitigate this, Android requires apps to obtain explicit user permission to access sensitive geolocation capabilities.

By processing network flows in our automated tests, we identified a probable leak of street-resolvable geolocation data occurring without user consent. In 15 out of the 111 auto-tested apps—all by the developer BabyBus—we observed the current Wi-Fi router's MAC address being transmitted to *tdcv3.talkingdata.net*. Wi-Fi router MAC addresses are trivially street-resolvable using public router lookup APIs offered by WiGLE and Google Maps Geolocation.

We verified this result by pinpointing our offices using the MAC address sent to the third-party firm. (Figure 4) We also successfully verified this capability using the home router MAC addresses observed in our manual testers' data.

BabyBus was informed of our findings weeks before submitting this paper. They have since responded that they removed the TalkingData library from their apps as a result of our notice. This indicates that our automated testing system could also be of value to children's apps developers seeking to audit the bundled libraries that their products use.

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

TABLE I
SELECTED APP CATEGORIES THAT CAN POTENTIALLY BE USED BY CHILDREN.

B. Lumen Dataset

Next, we mined anonymized traffic traces gathered from over 1,400 Lumen users for potential COPPA violations. As users interact with apps during the course of normal use, Lumen reports those apps' tracker activity to our database. For user privacy, Lumen only records the *presence* of unique identifiers in third-party traffic, not the values of the identifiers themselves. Lumen searches for instances of the device's own identifiers in outbound traffic.

We acknowledge that malicious developers may escape detection by arbitrarily encoding information at transmission—indeed, such methods strongly suggest a bad-faith attempt to hide objectionable activity from users that are sufficiently technologically savvy to examine network flows. This, however, is outside the scope of our COPPA compliance testing; our results are therefore a lower-bound of COPPA violations in the wild.

Our dataset contains 3,458 mobile apps run by these Lumen users. This dataset complements the artificial UI events generated by our testbed with traffic monitored *in situ*.

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"). 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 I.

Our analysis revealed 154 COPPA-relevant apps sharing unique identifiers with third-party services. 82 of them fall in the Educational or Education categories, followed by Casual (24 apps) [2] and Puzzle (20 apps). 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 do not need to request any specific permission [27]: 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, by bypassing the Android mechanisms that protect unique identifiers. 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.)

[2]This category comprises games such as Bubble Beach and My Talking Tom.

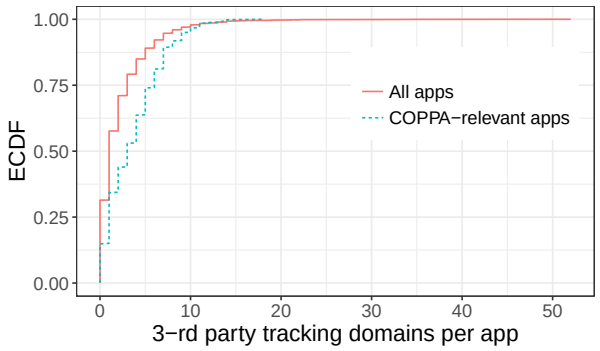

Fig. 5. CDF of the number of trackers for potentially COPPA-related apps (i.e., relevant ESRB categories such as Everyone, PEGI-3, PEGI-7 and PEGI-12 and app categories defined in Table I) ($N = 572$) and the total corpus of apps ($N = 3,458$).

We conclude our analysis with a comparison of the number of trackers found in apps in the categories listed in Table I 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 Lumen team [32]. As we can see in Figure 5, 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. Over 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 19 games reaching more than 10 third-party tracking and advertising domains. 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].

VII. 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, showing that not much has changed in five years [28], [29].

Acknowledgments

This research was supported by the United States Department of Homeland Security's Science and Technology Directorate under contract FA8750-16-C-0140, the Center for Long-Term Cybersecurity (CLTC) at UC Berkeley, the National Science Foundation under grants CNS-1318680 and CNS-1564329, the European Union under the H2020 TYPES (653449) project, and Data Transparency Lab Grant 2016 program. The content of this document does not necessarily reflect the position or the policy of the U.S. Government, European Union, or any other sponsor, and no official endorsement should be inferred.

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." [Online]. Available: https://www.ftc.gov/news-events/press-releases/2014/02/ftc-settles-childrens-gaming-company-falsely-claiming-comply
[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] "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.
[16] "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.
[17] Android Developer's Documentation, "Android developers: UI/application exerciser monkey," http://developer.android.com/tools/help/monkey.html.
[18] 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*, ser. OSDI'10. Berkeley, CA, USA: USENIX Association, 2010, pp. 1–6. [Online]. Available: http://dl.acm.org/citation.cfm?id=1924943.1924971
[19] European Commission, "Protection of personal data," http://ec.europa.eu/justice/data-protection/, November 24 2016.
[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*, ser. SOUPS '12. New York, NY, USA: ACM, 2012. [Online]. Available: http://doi.acm.org/10.1145/2335356.2335360
[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*. ACM, 2015, pp. 1111–1120.

[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. 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.

[27] 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.

[28] 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.

[29] ——, "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.

[30] ——, "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.

[31] 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.

[32] N. Vallina-Rodriguez, S. Sundaresan, A. Razaghpanah, R. Nithyanand, M. Allman, C. Kreibich, and P. Gill, "Tracking the trackers: Towards understanding the mobile advertising and tracking ecosystem," *arXiv preprint arXiv:1609.07190*, 2016.

[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)*. Washington, D.C.: USENIX Association, Aug. 2015, pp. 499–514. [Online]. Available: https://www.usenix.org/conference/usenixsecurity15/technical-sessions/presentation/wijesekera