



DIFFAUDIT: Auditing Privacy Practices of Online Services for Children and Adolescents

Olivia Figueira

University of California, Irvine
Irvine, California, USA
olivia.f@uci.edu

Athina Markopoulou

University of California, Irvine
Irvine, California, USA
athina@uci.edu

Rahmadi Trimananda

University of California, Irvine
Irvine, California, USA
rtrimana@uci.edu

Scott Jordan

University of California, Irvine
Irvine, California, USA
sjordan@uci.edu

Abstract

Children's and adolescents' online data privacy are regulated by laws such as the Children's Online Privacy Protection Act (COPPA) and the California Consumer Privacy Act (CCPA). Online services that are directed towards general audiences (*i.e.*, including children, adolescents, and adults) must comply with these laws. In this paper, first, we present DIFFAUDIT, a platform-agnostic privacy auditing methodology for general audience services. DIFFAUDIT performs differential analysis of network traffic data flows to compare data processing practices (i) between child, adolescent, and adult users and (ii) before and after consent is given and user age is disclosed. We also present a data type classification method that utilizes GPT-4 and our data type ontology based on COPPA and CCPA, allowing us to identify considerably more data types than prior work. Second, we apply DIFFAUDIT to a set of popular general audience mobile and web services and observe a rich set of behaviors extracted from over 440K outgoing requests, containing 3,968 unique data types we extracted and classified. We reveal problematic data processing practices prior to consent and age disclosure, lack of differentiation between age-specific data flows, inconsistent privacy policy disclosures, and sharing of linkable data with third parties, including advertising and tracking services.

CCS Concepts

- Security and privacy → Human and societal aspects of security and privacy;
- Social and professional topics → Government technology policy; Children; Adolescents; Privacy policies;
- Networks → Network measurement.

Keywords

Privacy auditing; COPPA; CCPA; online services for children and adolescents

ACM Reference Format:

Olivia Figueira, Rahmadi Trimananda, Athina Markopoulou, and Scott Jordan. 2024. DIFFAUDIT: Auditing Privacy Practices of Online Services for Children and Adolescents. In *Proceedings of the 2024 ACM Internet Measurement Conference (IMC '24), November 4–6, 2024, Madrid, Spain*. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3646547.3688416>



This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike International 4.0 License.

IMC '24, November 4–6, 2024, Madrid, Spain
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0592-2/24/11
<https://doi.org/10.1145/3646547.3688416>

Children and Adolescents. In *Proceedings of the 2024 ACM Internet Measurement Conference (IMC '24), November 4–6, 2024, Madrid, Spain*. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3646547.3688416>

1 Introduction

The Children's Online Privacy Protection Act (COPPA) in the United States (US) regulates online service providers' collection and sharing of personal information about children under the age of 13 [112]. With the increased presence of both children and adolescents online [11, 26], more recent privacy legislation such as the California Consumer Privacy Act (CCPA) [101] have incorporated additional privacy protections, akin to COPPA, for users under the age of 16. For example, CCPA forbids businesses from selling and sharing the personal information of users younger than 16, unless they have received consent. Additionally, online services, such as social media and gaming platforms, have been under increased scrutiny for their lack of safety and privacy for young users, prompting proposals for more stringent laws and regulations in the US [48, 57, 100, 104]. Methodologies for auditing data processing done by online services for young users can inform and enforce such regulations.

In this paper, we develop an auditing methodology to analyze the data collection and sharing practices of online services for young users and reveal potential violations of privacy laws. More specifically, we study online services directed towards general audiences, which include children, adolescents, and adults, and that allow users to make an account and disclose their age. We identified the following services out of the top 100 most popular apps on the Google Play Store: Roblox [95], Minecraft [73], TikTok [106], YouTube [36] and YouTube Kids [37], Duolingo [20], and Quizlet [91]. When creating an account on such online services, users are prompted to enter their age, which should impact the data collection and sharing that is done based on whether the user is a child (younger than 13), adolescent (between 13 and 16), or adult (older than 16), which are grouped in this way based on COPPA and CCPA. Also, upon creating an account and agreeing to the privacy policy and terms, consent to the services' disclosed data processing practices is given, either by the parent or by the user themselves depending on the user's age. Given the age-specific regulations in COPPA and CCPA, we expect well-behaved online services for general audiences to perform different data collection and sharing practices when users are logged in vs. out, as well as across different age groups.

While prior works have developed frameworks and tools for privacy auditing and measurement for several platforms, such as mobile, browsers, and virtual reality [4, 47, 92, 102], to the best of our knowledge, none have audited (i) general audience services (ii) in the context of both COPPA and CCPA for children and adolescents and (iii) their behaviors before and after providing consent and user age. Towards this goal, we make two contributions: (1) we design and implement the DIFFAUDIT auditing methodology, and (2) we apply it to a set of popular general audience services and report our findings.

First, we design **DIFFAUDIT** – a methodology for auditing the practices of general audience services and comparing data collection and sharing practices across different user age groups, *i.e.*, children, adolescents, and adults. DIFFAUDIT has the following characteristics:

- *Platform-Agnostic Differential Analysis:* DIFFAUDIT solely relies on the availability of network traffic collected when using a service on the end device (*e.g.*, a mobile app, a browser, etc.) for different age group profiles. We perform differential analysis to compare data flows (defined as *<data type category, destination>*) observed across age groups, as well as before and after consent is given and user age is disclosed. Destinations are also categorized as first/third parties and advertising and tracking services (ATS) using block lists.
- *Data Type Classification:* We develop a novel data type classification method using OpenAI’s GPT-4 [81] to classify data types extracted directly from network traffic into categories. We also develop a *data type ontology* rooted at the COPPA and CCPA definitions of identifiers and personal information [13, 16] (see Table 5 in Appendix). Our method outperforms previously used methods both in accuracy and coverage, which we discuss in Section 3.2.2. Prior works [1, 4, 9, 40, 47, 51, 92, 109, 120] focused on data types largely limited to identifiers, whereas our ontology also includes behavioral data types and is rooted in privacy laws rather than only extracted from the network traces. With our method, we were able to identify considerably more data types in the network traffic than prior work.

Second, we apply DIFFAUDIT to popular general audience services on both their mobile and website platforms. In the network traffic dataset we collected, we observed more than 440K outgoing requests, with 964 unique domains and 326 unique eSLDs contacted, and we extracted 3,968 unique data types from the packets and 5,508 unique data flows. In all the services we audited with DIFFAUDIT, we found practices for children and adolescent users that are concerning for COPPA and CCPA compliance, including the following:

- *Data Processing Before Consent:* All services collected users’ identifiers and personal information while logged out, and all but one of the services shared users’ identifiers and personal information with third party ATS while logged out. This raises concerns regarding compliance due to a lack of consent and knowledge of user age.
- *Differential Auditing:* All but one of the services engaged in data processing practices that were not disclosed in their privacy policy, including while users are logged in and out,

which is problematic for transparency and compliance. All general audience services we studied engaged in similar data processing practices for both young and adult users.

- *Platform Differences:* We observed some data flows that are unique to the mobile app and website platforms. We found that the observed data flows unique to the mobile apps were all related to sharing data with third parties, including personal and device identifiers, geolocation, and app usage behaviors. This reveals some potential internal differences in services’ mobile vs. website platforms.
- *Data Linkability:* All but one of the services shared linkable data (*i.e.*, identifiers and personal information shared with the same third parties) with both ATS and non-ATS third party domains, both before and after consent. These services also sent the linkable data to similar destination domains, without much differentiation between age groups.

We make DIFFAUDIT’s implementation and datasets available at <https://athinagroup.eng.uci.edu/projects/childrens-privacy/>.

The rest of the paper is structured as follows. Section 2 provides background and discussion of related work. Section 3 presents the DIFFAUDIT methodology, including: the network traffic collection (Section 3.1), the construction of data flows, and the data type classification method (Section 3.2). Section 4 presents the results of applying DIFFAUDIT to popular online services. Section 5 discusses implications of our findings, recommendations for service providers and regulators, and ethics considerations. Section 6 concludes the paper. The Appendices include details regarding our data type classification method and data type ontology (Table 5).

2 Background & Related Work

Section 2.1 discusses how COPPA and CCPA motivate the design of DIFFAUDIT and Section 2.2 discusses general audience services that we investigate. Sections 2.3 and 2.4 discuss related work on privacy auditing and children’s and adolescents’ privacy, respectively.

2.1 COPPA & CCPA

COPPA and CCPA regulations for users under 16 motivate the design of our auditing framework. COPPA applies to services that are directed towards children, which includes services directed towards general audiences composed of users of all ages, including children, adolescents, and adults. According to COPPA, a “Web site or online service directed to children means a commercial Web site or online service, or portion thereof, that is targeted to children”¹, which may be judged based on features and content of the service (*i.e.*, “subject matter, visual content, use of animated characters or child-oriented activities and incentives, music or other audio content”²) and knowledge of children using the service [29]. COPPA prohibits such online service providers from collecting any personal information prior to knowing any given user’s age, since the audience could contain children under 13, and even once they know the user’s age, they require opt-in consent from the child’s parent or guardian to collect and share such personal information, except for activities for the support of internal operations [112].

¹16 C.F.R. § 312.2 “Web site or online service directed to children”

²16 C.F.R. § 312.2 “Web site or online service directed to children” (1)

Under CCPA, businesses are not allowed to sell nor share users' personal information for users younger than 16 years old without the user's prior affirmative authorization (*i.e.*, opt-in consent). For users under 13 years old, businesses are not allowed to sell nor share users' personal information without their parent's or guardian's opt-in consent, similar to COPPA [12]. Additionally, CCPA states, "A business that willfully disregards the consumer's age shall be deemed to have had actual knowledge of the consumer's age"³. While some of these terms (*i.e.*, willfully disregard, actual knowledge) are not explicitly defined, this statement can be interpreted to mean that online services should not be selling nor sharing users' data at all until they determine their age and obtain opt-in consent. Note that we do not audit the adult data flows against CCPA in this work, rather we only audit the child and adolescent data flows and use the adult data flows for comparison.

General audience services should wait until they determine the user's age and obtain consent, either from the user or their parent depending on their age, to perform the appropriate data collection/sharing under COPPA and CCPA. Thus, we develop an auditing methodology, DIFFAUDIT, for general audience services based on network traffic to both investigate whether user data is being collected/shared prior to consent and age disclosure (*i.e.*, while logged out) and how similar the data collection/sharing practices are for different user age groups (*e.g.*, data collected/shared for child vs. adult and adolescent vs. adult accounts). Moving forward, we will use the term "logged out" to refer to the state prior to user consent and age disclosure.

We also defined a data type ontology using the definitions of identifiers and personal information from COPPA and CCPA [13, 16]. See a description in Section 3.2.2 and details in Table 5 in Appendix A.

2.2 General Audience Services

With our auditing goals in mind, we establish the following criteria for the types of services we wish to audit with this framework: (i) the services should target general audiences, *i.e.*, including children, adolescents, and adults, and (ii) the services should prompt users to make an account so that we can disclose the user's age and provide consent to the behaviors disclosed in the privacy policy. To select services to audit, we searched through the top-100 most popular games and apps on the Google Play Store and manually inspected each app's privacy policy to determine the target audience and whether the app fit our criteria. At the time of this study, the following were the only services that met these criteria: Duolingo [20], Minecraft [73] (owned by Microsoft), Quizlet [91], Roblox [95], TikTok [106], and YouTube [36], which also includes YouTube Kids [37]. Roblox and Minecraft are gaming platforms, TikTok and YouTube are social media and content platforms, and Duolingo and Quizlet are educational services. We investigate each service both on the mobile and the web platforms, in addition to desktop apps for Roblox and Minecraft.

These are all important services: the websites of these six services are among the most popular on the top 1M Tranco list at the time this work was conducted (fall 2023), and Roblox, TikTok, and

YouTube were among the top 100 [61, 113]. Furthermore, the corresponding apps for these services have been cumulatively downloaded more than 12 billion times on the Google Play Store [35] and have received close to 280 million reviews around fall 2023.

All selected services are directed towards children, adolescents, and adults according to their terms of service, and thus both COPPA and CCPA apply. For example, YouTube has both the regular YouTube service, which we audit based on CCPA since it allows adolescents, and YouTube Kids, which we audit based on both CCPA and COPPA. The selected services allow users of all ages to access their services, with some services requiring active parental consent for users under 13, such as requiring a parent's email in the account creation process.

2.3 Online Privacy Analysis & Auditing

Prior works have developed mechanisms for analyzing data collection and sharing practices across diverse platforms through static, dynamic, or network traffic analysis and have audited such practices against privacy policies and laws. However, few have focused specifically on children's and adolescents' privacy in comparison. There exist many tools for privacy analysis and auditing for mobile devices and apps that are based on network traffic analysis to investigate data collection and sharing practices as well as presence of ATS [1, 4–6, 40, 51, 120]. Others have used dynamic taint analysis to monitor real-time data sharing on Android [114] as well as static analysis to analyze third-party tracking on the Google Play Store [93]. For websites and browsers, prior works have developed privacy measurement tools for network traffic and user advertising consent preferences [102, 122]. Researchers have also developed frameworks to investigate privacy of emerging platforms, such as virtual reality device app behaviors [92], personal information exposure and prevalence of ATS in smart TVs [43, 50], and smart speaker tracking behaviors [111]. Additionally, prior works have studied consent in third-party tracking and data sharing in Android apps [59, 109].

Specifically related to children's privacy, researchers have analyzed privacy practices of mobile apps intended for children [47], risky content for children in voice assistants [110], and targeted advertising towards children online [108]. Others have audited the privacy and security of parental control applications [103, 123], algorithmic personalization on TikTok for young users [70], and IoT smart toys for children [38]. While these works revealed privacy issues and advanced transparency and privacy, more investigation is needed as technology continues to evolve and become increasingly accessible to children and adolescents, introducing new risks.

The closest to our work is that of Reyes et al. [47], which analyzed data collection and sharing practices of Android mobile apps specifically for children w.r.t. COPPA using dynamic analysis. Our work differs in that (i) we specifically investigate general audience services, which have not been previously studied, (ii) on both mobile and website platforms, and (iii) we apply a differential analysis approach based on age and consent status (*i.e.*, while logged in and out) in context of both COPPA and CCPA. In addition, our data type classification method expands beyond prior work and allows us to identify considerably more data types.

³CAL. CIV. Code § 1798.120(c)

2.4 Children's & Adolescents' Privacy

The research community has also investigated how users, including children, adolescents, and parents, approach children's and adolescents' privacy from different perspectives.

Prior works have studied how young users and parents understand and make decisions about online privacy, such as privacy conceptualizations across technologies in general [55, 56, 62, 88], virtual reality systems [23], social media [17, 18, 46, 76], and smart toys [27, 77]. Researchers have also investigated advertising directed towards children and adolescents through content analysis and studies of user perceptions [14, 69, 117, 118]. Also, researchers have studied the relationship between parents and their children regarding technology usage and privacy management, including experiences with parental controls [8, 85] and parental oversight strategies [67, 68, 83, 86]. Prior works have also studied age assurance and parental consent mechanisms from legislative and policy-making perspectives [2, 97, 99], through literature reviews [15], or via small-scale analyses of mechanisms currently in use on popular platforms [65].

Prior works have also developed tools to both protect user privacy and increase transparency, such as privacy protections for mobile apps [10, 30], smart toys [98], and facial recognition technologies [3, 84], as well as design insights for privacy labels and warnings [52, 71]. The research community has also studied barriers faced by developers in designing applications for children [7, 22] and has proposed design guidelines for children's privacy [31, 32, 64, 87].

We contribute to this body of work by demonstrating that services meant for children and adolescents continue to behave in privacy-pervasive ways. Additionally, many of the observed behaviors are difficult to protect against with existing tools for children's and adolescents' privacy.

3 The DIFFAUDIT Methodology

In this section, we describe how we collect and analyze the outgoing network traffic from the mobile apps and websites while using the services of interest. In Section 3.1, we describe the setup and process for network traffic collection, including collecting traffic while logged out and for each user age group while logged in. In Section 3.2, we describe how we post-process the network traffic traces to extract the data flows, including our data type classification method to categorize extracted data types and how we categorize the packet destinations.

We make DIFFAUDIT's implementation and datasets available at <https://athinagroup.eng.uci.edu/projects/childrens-privacy/>.

3.1 Network Traffic Collection

Our auditing framework, as shown in Figure 1, relies on network traffic analysis. To be able to compare network traffic transmitted on each service for different age categories, we created three profiles on each service, one for each of the three age categories: child (younger than 13), adolescent (at least 13 and younger than 16), and adult (16 and older). We collected network traffic for each profile on the mobile app and website for each service, amounting to three traces per app and per website. In addition, we collected three types of network traces to compare the data processing practices

before, during, and after account creation: (1) *Account creation trace*: We collected data throughout the entire process of creating an account (which includes providing user age and consenting to the privacy policy and terms of service). (2) *Logged-in trace*: We collected data only during service usage, meaning we logged in to each account and then began collecting network traffic only while logged in (provides user age due to being logged in). (3) *Logged-out trace*: We collected data only while logged out, so there are no age-specific traces in this category (*i.e.*, consent and age have not been given/disclosed). We examine the default state of these services, meaning we did not opt in nor out of any data collection/sharing for any of the age categories beyond what was consented to in the privacy policies.

While there exist tools for automated interaction, we chose to use manual interaction because it enables exhaustive exploration of the services and all their features, allowing for more in-depth exploration of each service. Apps and websites often change their layouts and services depending on the age of the user currently signed in, such as having different features for child versus adult accounts, which complicates automated interaction methods. Manual interaction also resembles realistic service usage and has been used in prior work for privacy auditing tasks for this reason [40, 50, 51, 92]⁴.

For each kind of trace, we attempted to utilize every feature available in the service and exhaust all possible interactions. The account creation and logged-in traces should be a minimum of five minutes, whereas the logged-out trace category is usually shorter as we are limited in functionality while logged out. Thus, we have about 10-15 minutes of data collection time per age category for each service on both platforms (*i.e.*, about one hour total per service). To maintain consistency in our network traffic collection, we followed the same order of operations for each website and app: (1) logged-out trace, (2) account creation trace, and (3) logged-in trace. For each trace, we also followed the same interaction steps within the app or website when possible (*i.e.*, some features and layouts change based on the age group). For example, on TikTok, we would click on all buttons and menus in the same order and interact with the content feed in the same manner across each age profile to ensure we interacted with everything possible in a consistent fashion. Table 1 presents the summary of our resulting network traffic dataset per service, with mobile and website data merged. We observed 964 unique domains and 326 unique eSLDs among 440,513 outgoing packets in total.

3.1.1 Mobile Apps. To perform the network traffic collection on mobile apps, we used PCAPdroid [24] on a Pixel 6 Android mobile device. PCAPdroid is an open source app that can track, analyze, and block the connections of other apps running on the same Android device. PCAPdroid simulates a VPN to capture the network traffic without root: it does not use a remote VPN server and data is processed locally on the device. However, to decrypt the network traffic, we had to root the device and install PCAPdroid's user certificate as a trusted certificate [25]. To collect network traffic, we launch the target app, begin the trace within PCAPdroid, and use

⁴For the six popular general audience services we study in depth in this paper, manual interaction is possible and preferred for the reasons mentioned above. However, even if the interaction with the services becomes automated in the future, the rest of the DIFFAUDIT framework still holds.

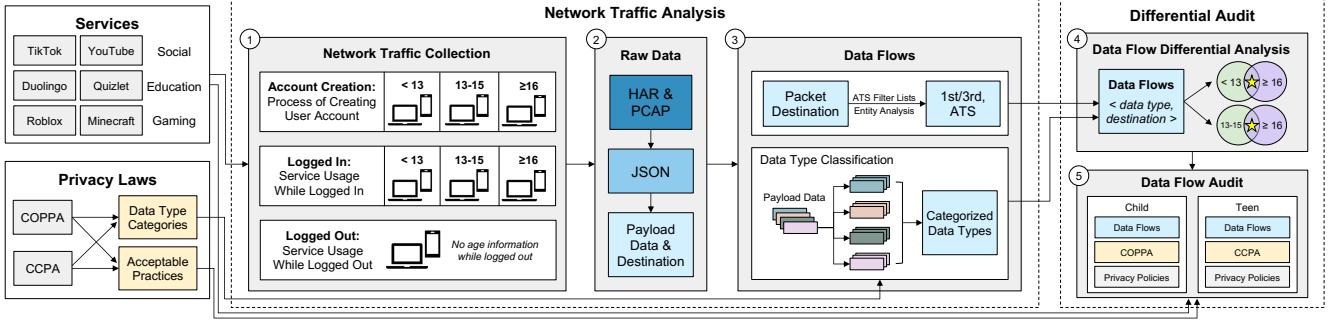


Figure 1: DIFFAUDIT Framework Overview. First, we select the general audience services for auditing. Next, we perform **network traffic analysis**: (1) Collect network traffic in three different ways: **account creation trace** (we collect traffic while creating a user account), **logged-in trace** (we collect traffic only while logged in to each account on the website/app and throughout using the service), and **logged-out trace** (we collect traffic only while logged out of any account on the website/app). (2) For each trace, we convert the raw HAR (for website) or PCAP (for mobile) data to JSON and extract the payload data and destinations of the packets. (3) We construct the data flows by processing packet destinations (i.e., 1st or 3rd party and entity analysis) and perform data type classification using GPT-4 classifier and COPPA/CCPA data type ontology. Next, we perform the **differential audit**: (4) compare the data flows by age group and (5) audit the flows in context of each age group, applicable law, and information from each service’s privacy policy.

Table 1: Network Traffic Dataset Summary. The same domains and eSLDs may appear across the different services, thus their totals are based on their unique totals. YouTube includes YouTube Kids.

Service	Domains	eSLDs	Packets	TCP Flows
Duolingo	122	69	60,909	1,466
Minecraft	136	56	134,852	2,004
Quizlet	532	257	88,102	6,158
Roblox	152	24	103,642	2,302
TikTok	80	14	32,234	2,412
YouTube	76	15	20,774	226
Total	964	326	440,513	14,568

the target app to generate traffic. When we are done with the trace, we stop the collection in PCAPdroid and save the generated PCAP file to local device storage. PCAPdroid also generates a TLS key log file, which we save to local device storage and use to decrypt the generated PCAP file with Wireshark [116]. We clear the storage and cache of the app in the device settings between traces. We attempted to get as much decrypted traffic as possible by also employing certificate pinning bypass techniques using Frida [78] along with PCAPdroid. We include all collected traffic, both encrypted and decrypted, in our analysis.

3.1.2 Websites. To perform the network traffic collection on the website services, we used the Chrome browser [34] and Chrome DevTools Network Panel [58]. We ensured the browser was reset to its default settings and cleared all browsing data (i.e., all browsing history, cookies, site data, and cached data) between each trace. For each trace, we ensured that the “Preserve logs” checkbox is selected before recording to ensure all traffic is collected when new sites are opened during the same session. After completing each network trace, we export the trace as a HAR (HTTP Archive) file [49], as provided by the Network Panel.

3.1.3 Desktop Apps. Roblox and Minecraft also provide complementary desktop apps with more features, namely the gaming features of these services. Thus, we also collected network traffic in the same manner from the desktop app versions of these services using Proxyman [89], which is a man-in-the-middle system that we use to capture and decrypt traffic through SSL proxying, and we export the traces to HAR files, similarly to the website process.

3.2 Network Traffic Post-Processing

We use PCAPdroid to generate a PCAP file and TLS key log file for each mobile app trace. We use the Wireshark functionality editcap [115] to embed the TLS keys into the PCAP file and generate a new decrypted version of the PCAP file, which we then use for the subsequent network traffic post-processing steps. We convert the HAR and PCAP files to JSON format and extract all outgoing requests from each trace, as we are interested in the data leaving our device.

3.2.1 Data Flows in Context. Starting from a raw, clear-text network traffic trace, we extract each packet’s destination and data types from the packet headers and payloads, respectively. We use that information to construct *data flows*, defined as a pair $\langle \text{data type category}, \text{destination} \rangle$, which is consistent with the terminology in [9, 39, 92]. We determine the appropriateness of a data flow based on the user’s age and logged-in/out status (i.e., indicating consent) in context with COPPA and CCPA. This can be thought of as a special case of appropriate information flows in the contextual integrity framework [41, 119].

3.2.2 Data Type Classification Methodology. We extract key-value pairs from the JSON-structured data in the packets, and the keys serve as the raw data types, which is common in prior work [4, 6, 9, 40, 51, 92, 120]. To categorize these raw data types, prior works have utilized manual inspection [92], string-matching techniques using pre-determined regular expressions [4, 6, 120], supervised

Table 2: Data Type Categories From Our Ontology. Our proposed data type ontology is based on COPPA/CCPA. Data types observed in our dataset are marked with “*”.

Identifiers [13, 16]	
Name*	Customer Numbers
Linked Personal Ids.	Login Info*
Contact Info*	Device Hardware Ids*
Aliases*	Device Software Ids*
Reasonably Linkable Personal Ids.*	Device Info*
Personal Information [13, 16]	
Race	Precise Geolocation
Age*	Coarse Geolocation*
Language*	Location Time*
Religion	Communications
Gender/Sex*	Contacts
Marital Status	Internet Activity
Military/Veteran Status	Network Connection Info*
Medical Conditions	Sensor Data
Genetic Info	Products & Advertising*
Disabilities	App/Service Usage*
Biometric Info	Account Settings*
Personal History	Service Info*
Inference About Users*	

learning with natural language processing techniques [51], or a combination of these techniques [40]. These prior works utilize data type ontologies, developed heuristically, and often require prior knowledge of user identifiers to directly search for in the network traffic or a precise list of strings to match for with various permutations.

In contrast, we propose a *data type ontology*, which we define based on the COPPA and CCPA legal definitions of identifiers and personal information [13, 16]. The CCPA defines personal information as information that is reasonably linkable to a consumer or household, and it specifies a non-exhaustive list of personal information, such as name, postal address, unique personal identifier, IP address, characteristics of protected classifications under California or federal law, biometric information, electronic network activity information (e.g., browsing history, search history), geolocation data, and consumer behaviors and preferences [13]. We organized this list into categories, supplemented it with a list of personal information specified in COPPA, and added other personal and device identifiers and relevant information commonly observed in network traffic, namely (i) device information, which can enable device fingerprinting when combined with other attributes [60, 75], (ii) developer-oriented service information related to the implementation of the websites and apps (e.g., information about SDKs and APIs, website/app components and content placement), and (iii) network connection information (e.g., status of connection, protocol information), as listed in Table 2. Although these categories are not explicitly listed under the CCPA and COPPA, they fall within their definitions of personal information. Specifically, (i) device information can be used to create uniquely identifiable device fingerprints and thus is included in our ontology under “Device Identifiers”,

(ii) developer-oriented service information can reveal information about users’ behaviors based on the content or components rendered in a website/app and thus is included in our ontology under “User Interests and Behavior”, and (iii) network connection information can reveal data about users’ network communications and thus is included in our ontology under “User Communications”. Please see details regarding the full ontology in Table 5, Appendix A. Additionally, under the CCPA and COPPA, personal information does not include publicly available information and de-identified information, and therefore these kinds of data types are not included in the ontology.

We show an excerpt of the multi-level ontology in Figure 2 and the full ontology can be found in Appendix A in Table 5. Our ontology has four levels, where level 1 consists of these two main categories (identifiers and personal information). Level 2 splits into 8 categories (i.e., personal identifiers, device identifiers, personal characteristics, personal history, geolocation, user communications, sensors, and user interests and behaviors), which split into 35 level-3 categories, as shown in Table 2. Table 2 also indicates with an “*” which data types were observed in our dataset. Level 4 contains the data types belonging to each of the 35 level-3 categories. Levels 3 and 4 are used in our data type classification methodology, as we discuss later in this section.

Prior works that conduct data type categorization primarily focus on identifiers and include a limited number of data types [1, 4, 9, 40, 47, 51, 92, 109, 120], with [51] including more than the rest, but still only covers 12 of the 35 data types we include (i.e., device hardware/software identifiers, name, alias, gender/sex, age, contact information, contacts, marital status, precise/coarse geolocation, and login information). Out of the 35 data types we include, we identified 19 in our network traffic dataset. With our data type ontology, we are able to capture considerably more data types, including behavioral data.

Our *data type classification method* is built using our data type ontology and OpenAI’s GPT-4 8K context model [81] with their Chat Completion API [79], which can take in prompts and generate textual responses. Note that this work was completed prior to OpenAI’s launch of GPT-4 Turbo and other updates to their APIs⁵. Our approach, depicted in Figure 2, is similar to a few-shot classification framework, where we treat GPT-4 as a classifier that takes in the category labels (level-3 categories from our ontology) and data type examples in each category (level-4 data types from our ontology), and the model must classify our raw input texts using these categories.

A challenge with using GPT-4 is designing the input prompt, since the prompt is what guides the model in understanding and completing the given task, and it can alter the output depending on how the task is worded and what information is given. We experimented with several iterations of prompts and followed OpenAI’s GPT best practices guide [82]. We attempted to design a prompt that resulted in a reliable output format, which we required to be able to operationalize the output of the model for classification. We also had to determine what value to use for GPT-4’s temperature parameter, which dictates how creative the model will be in

⁵OpenAI DevDay, November 2023: <https://openai.com/blog/new-models-and-developer-products-announced-at-devday>

Table 3: GPT-4 Classification Model Sample Validation Results. Accuracy and coverage results are based on a random sample. The table shows total sample accuracy and per-confidence score results for each GPT-4 model varying by temperature and method.

Temperature or Method	Accuracy	Confidence 0.7		Confidence 0.8		Confidence 0.9	
		Accuracy	Labeled	Accuracy	Labeled	Accuracy	Labeled
0	0.72	0.76	363	0.82	292	0.87	221
0.25	0.73	0.77	374	0.83	294	0.84	230
0.5	0.69	0.72	368	0.78	297	0.80	224
0.75	0.66	0.71	353	0.78	271	0.81	187
1.0	0.65	0.70	356	0.74	286	0.78	201
Majority-Max	0.75	0.77	381	0.82	354	0.85	299
Majority-Avg	0.75	0.81	358	0.87	274	0.91	162

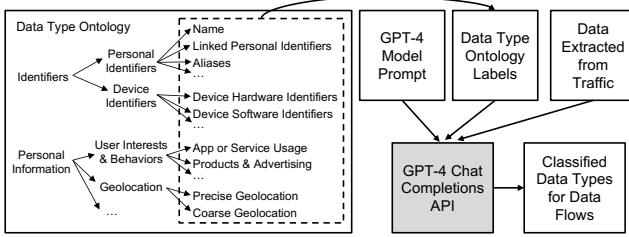


Figure 2: Data Type Classification System Diagram. On the left we present an excerpt of the *data type ontology*. The data type category labels (third level) are used as labels for the classifier. We input the *GPT-4 model prompt*, *data type ontology labels*, and *raw data types* extracted from the network traffic to the *GPT-4 Chat Completions API*, which outputs classification results.

its response and ranges from 0 to 2 [80]. For values greater than 1, we observed hallucinatory responses, so we only used values between 0 and 1 in increments of 0.25. Additionally, we prompted GPT-4 to output a confidence level along with each output to serve as a confidence threshold for our final labels. Due to the nature of network traffic data, there are input texts that will be unclear to the model, as well as to humans doing manual labeling, such as unknown acronyms or seemingly random characters that have internal meaning known only to the app developers, and we do not want to include such low confidence guesses in the dataset. See Appendix B for details about the model prompt.

Considering the inherent nondeterminism of GPT-4, we build a majority-vote model where we take the majority label assigned across all the different temperature models and assign that majority label as the final label to attempt to balance model creativity, accuracy, and nondeterminism. For the majority-vote model confidence score threshold, we either compute the accuracy and number of inputs labeled based on the maximum confidence score among the models that assigned the majority label, or we can use the average confidence score of the scores given by the models that assigned the majority label.

We also experimented with other labeling methods, namely fuzzy string matching with TF-IDF and BERT embeddings [66] and both zero- and few-shot classification with transformers [45, 63], but such approaches achieved poor accuracy when validated on a

manually-labeled sample of our dataset (see next section). Specifically, we utilized the library PolyFuzz [66] for fuzzy string matching, which enables customizable fuzzy string matching with Hugging Face transformers as well as word embedding methods such as TF-IDF and BERT. For the zero-shot classification approach, we used the zero-shot classification pipeline [45] provided by Hugging Face with bart-large-mnli [44], which is based on the BART model [74]. Finally, for the few-shot classification approach, we utilized SetFit [63], which is a few-shot learning framework. Overall, the raw network traffic includes a myriad of string formats, such as those that directly relate to their meaning (e.g., “email”, “username”), acronyms and abbreviations of well-defined terms (e.g., “os”, “rtt”), and well-defined terms concatenated with other text and/or punctuation (e.g., “pers_ad_show_third_part_measurement”, “IsOptOutEmailShown”). GPT-4 is better equipped to extract the context and meaning compared to previously used methods and thus achieves greater accuracy. Additionally, we do not require prior knowledge of user identifiers, allowing for a much more flexible approach.

To validate and parameterize our models, we require labeled data to serve as the ground truth so we can analyze the performance. Our dataset is unlabeled as it is composed of keys from the key-value pairs in the network traffic, and the task is to classify them using our ontology data type categories. Thus, two researchers independently and manually labeled a random sample of 10% (n=397) of the entire dataset of extracted network traffic data types using our classification categories. The two researchers originally had 84% agreement in their labeling, but came to a consensus after discussion and merging of their labels. Thus, the resulting labeled sample data was used as the ground truth for model validation. We ran each of the classification approaches and computed the accuracy based on this manually-labeled sample. Besides GPT-4, the best performing alternative was fuzzy string matching with TF-IDF embeddings at 31%, and the rest achieved even lower sample accuracy scores (zero-shot at 4%, few-shot at 16%, and fuzzy string matching with BERT at 18%). We present in detail the sample accuracy results and coverage counts at different confidence levels for each GPT-4 model in Table 3. The GPT-4 models perform significantly better than the other classification approaches mentioned, and the majority-vote approaches perform better than the individual temperature GPT-4 models when taking into account the confidence thresholds. To maximize accuracy while maintaining reasonable coverage, we use the labels resulting from the majority-average approach at the 0.8

confidence threshold in our final labeling scheme, resulting in 2,807 out of the 3,968 (70%) keys from the traffic included in the final labeled dataset.

Overall, this data type classification approach offers flexibility in its design through the confidence thresholds and options regarding single temperature-based models or majority-vote models, depending on the users' preferences w.r.t. coverage, accuracy, and cost. Additionally, our method produces a set of labeled network traffic payload data that can be used to train smaller models that can be run locally instead.

3.2.3 Destinations. To complete the data flows, we also post-process the packet destinations. To determine the fully qualified domain name (FQDN) destination for each request, we extracted the FQDN from the destination URL in each packet in the HAR and PCAP traces. Then, we extracted the effective second level domain (eSLD) using the tldextract library [53]. Based on the eSLD, we determine the parent organization owner of this domain, using whois [94] and the DuckDuckGo Tracker Radar dataset [19] if possible. The Tracker Radar dataset consists of the most commonly contacted third-party domains and has information about ownership, classification, and behavior, such as fingerprinting.

We categorize the domains as first parties if we can match it to the name of the service (*i.e.*, roblox.com) or if we identify that the parent owner of the domain is the same as that of the service. Otherwise, we categorize the domain as a third party. Furthermore, we identify which domains are likely to be advertising and tracking services (ATS) using ATS block lists [105] based on the FQDN, as prior works have also done [9, 50, 92, 111]. If any of the block lists results in a block decision for a particular domain, we label that domain as an ATS and store that information in a separate label, which we use in our data flow analysis. Overall, we identify whether each domain is a first party, a first party labeled as an ATS, a third party, or a third party labeled as an ATS.

4 Results

In this section, we analyze the results of our audits and discuss the findings. Section 4.1 discusses data flows observed for each service, including for the logged-out traces and per-age data flows, and compares the behaviors against each service's privacy policy. Section 4.2 presents the data linkability results.

4.1 Differential Data Flow Auditing

Table 4 shows the data flows observed across all network traces per service, with account creation and logged-in traces merged together. "Collect" indicates data sent to first parties, and "share" indicates data sent to third parties. The table indicates with different symbols whether data flows were observed on web, mobile, both, or neither, as shown in the caption. Recall that we were not able to collect all the network traffic in clear-text on the mobile apps, so the seeming lack of data flows unique to mobile is likely due to this difference.

Since there are 19 data type categories observed out of the 35 possible categories, as shown in Table 2, we only show the abstracted categories (level 2 from ontology) due to space, resulting in six data type categories shown in Table 4. For the purposes of auditing against COPPA and CCPA, the six abstracted categories suffice as

they represent the categories that are covered in the laws, as shown in the full ontology in Table 5 in Appendix A.

We analyze whether there are differences between the age-specific and the logged-out traces to determine whether the data processing is changing as we expect and whether it is appropriate for these categories. For the child category, if we observe any data flows among these law-protected categories either being collected or being shared (*i.e.*, all the collect and share columns in Table 4) that are not clearly disclosed in the privacy policy, those data flows would raise concerns and necessitate further investigation, such as by regulators. For the adolescent trace category, if we observe any data flows among these law-protected categories being shared with third parties, particularly ATS domains that indicate non-functional data flows, without clear disclosure in the privacy policies, those data flows would raise concerns. Next, we will discuss the logged-out (pre-consent) category data flows, auditing results for each individual service, and data flow differences observed between the services hosted on different platforms (*i.e.*, web, mobile, or both).

4.1.1 Logged-Out Traces. Across Table 4, we observe many data flows in the logged-out category across all services. All the data type categories except for geolocation were observed being collected by first party domains while logged-out across all services (only Roblox did not collect geolocation). Additionally, many of the other data types were either collected by first party ATS, shared with third parties, or shared with third party ATS. For example, we observed device identifiers and user communications data being shared with third party ATS for all services except YouTube. Additionally, we observed data flows containing all data type categories being shared with third party ATS for Roblox, Duolingo, and Quizlet.

It is concerning that personal information and identifiers are collected and/or shared prior to the user informing the service of their age and prior to giving consent through the account creation procedure, particularly for the child and adolescent users for whom services must obtain opt-in consent. General audience services should wait until they are informed of user age and receive the appropriate consent so that they proceed with the appropriate data collection and sharing for the age category.

Key Takeaways: All of the services engaged in data collection and/or sharing prior to consent and age disclosure, which is problematic for child and adolescent users for whom these services must obtain opt-in consent. All but one of the services (YouTube) was observed sharing identifiers and personal information with third party ATS while logged out.

4.1.2 Age-Specific Traces & Privacy Policy Analysis. Next, we will discuss the data flows observed for each individual service under audit based on Table 4 in addition to details about what data flows were disclosed in each service's privacy policy at the time of data collection (fall 2023).

Duolingo: We observed that all the listed data type categories, including personal and device identifiers, personal characteristics, geolocation, user communications, and user interests and behaviors, were collected by first parties and shared with third party ATS across all age categories. This is alarming both regarding the lack of difference between age groups and the data sharing particularly to third party ATS in the child and adolescent traces.

Table 4: Data Flows Observed by Age Category for Website and Mobile Platforms. Data flows are indicated by the presence of a symbol in the grid between a data type category and destination. The age-specific flows are merged from the account creation and logged-in traces. “Collect” indicates data sent to first parties. “Share” indicates data sent to third parties.

Service	Data Type	Child	Adolescent	Adult	Logged Out
		Collect 1st Collect 1st ATS Share 3rd	Collect 1st ATS Share 3rd ATS Collect 1st	Collect 1st ATS Share 3rd ATS Collect 1st	Collect 1st ATS Share 3rd ATS Collect 1st
Duolingo	Personal Identifiers	•	—	→	•
	Device Identifiers	•	—	•	•
	Personal Characteristics	•	—	→	•
	Geolocation	•	—	—	•
	User Communications	•	—	•	•
	User Interests and Behaviors	•	—	•	•
Minecraft	Personal Identifiers	•	•	□	—
	Device Identifiers	•	•	•	•
	Personal Characteristics	•	•	•	•
	Geolocation	•	→	→	□
	User Communications	•	•	•	•
	User Interests and Behaviors	•	•	→	•
Quizlet	Personal Identifiers	•	—	•	→
	Device Identifiers	•	—	•	•
	Personal Characteristics	•	—	•	•
	Geolocation	→	—	•	—
	User Communications	•	—	•	•
	User Interests and Behaviors	•	—	•	•
Roblox	Personal Identifiers	•	•	□	→
	Device Identifiers	•	•	•	•
	Personal Characteristics	•	•	•	•
	Geolocation	→	—	—	→
	User Communications	•	•	•	•
	User Interests and Behaviors	•	•	•	•
TikTok	Personal Identifiers	→	→	—	—
	Device Identifiers	•	•	→	□
	Personal Characteristics	→	→	—	—
	Geolocation	→	—	—	—
	User Communications	•	•	→	—
	User Interests and Behaviors	→	→	—	—
YouTube	Personal Identifiers	→	—	—	—
	Device Identifiers	→	—	—	—
	Personal Characteristics	→	—	—	—
	Geolocation	→	—	—	—
	User Communications	→	—	—	—
	User Interests and Behaviors	→	—	—	—

• : data flow observed on both website and mobile

→ : data flow observed only on website platform

— : data flow not observed on either platform

□ : data flow observed only on mobile platform

According to Duolingo’s privacy policy, for users under 16 “advertisements are set to non-personalised” and “third-party behavioral tracking is disabled” [21]. However, we observed several data types being sent to third party ATS for both the child and adolescent groups that suggest personalized ads and behavioral tracking, including both personal and device identifiers, personal characteristics, geolocation, and user interests and behaviors. We question the purpose of these data flows to third party ATS if not for behavioral tracking and personalized ads. Also, the same data type categories

were shared with third party ATS for the adult trace, which further casts doubt upon the purpose of these data flows.

Minecraft: For all the age-related traces, we observed that all listed data type categories were collected by first parties (ATS and non-ATS) and shared with non-ATS third parties. For the child and adolescent traces, we observed all listed data type categories except personal identifiers being shared with third party ATS, which means device identifiers, personal characteristics, geolocation, user communications, and user interests and behaviors about users under 16 were observed being shared with third party ATS. The only

difference between the adult and child/adolescent traces is that personal identifiers were not shared with third party ATS for the latter, otherwise they contain the same categories of data flows.

The privacy policy provided by Microsoft, which owns Minecraft, stated the following: “we do not deliver personalized advertising to children whose birthdate in their Microsoft account identifies them as under 18 years of age” [72]. Since we observed data flows with identifiers and personal information shared with third party ATS in both the child and adolescent traces, we question the consistency of these behaviors and the privacy policy disclosure, since ATS are known to be related to personalized advertising. Additionally, it is interesting that we observed sharing of all listed data type categories except personal identifiers with third party ATS for users under 16, but for the adult trace, personal identifiers were included. We are skeptical that simply removing the personal identifiers would result in de-identified data under CCPA, considering the other data types shared, and we question whether this was the intention.

Quizlet: We observed that all listed data type categories were collected by first parties, shared with third parties, and shared with third party ATS for all age traces. The lack of differentiation between the age traces and these collection and sharing behaviors overall are highly concerning for both the child and adolescent traces. We also observe the same behaviors for the logged-out trace, including all data type categories being shared with third party ATS, which raises concerns regarding their privacy policy disclosures.

The Quizlet privacy policy states the following for child users: “We may use aggregated or de-identified information about children for research, analysis, marketing and other commercial purposes” [90]. Thus, when consent is given for child users, they may share some personal information with third party ATS, which we have observed. However, we observed the same data type categories being shared with third party ATS while the user is logged out, which is problematic under COPPA and CCPA. While they may assert these data are de-identified and shared in aggregate, we observed all data type categories being shared and are skeptical that they are de-identified.

Roblox: The data flows for Roblox were largely the same across traces. The child, adolescent, and adult traces revealed all six listed data type categories being collected by first party non-ATS and shared with third party ATS. Additionally, all of the categories except geolocation were shared with third parties for all age-related traces. The only difference for the logged-out category compared to the age-related traces is that personal identifiers were not shared with non-ATS third parties and geolocation was not collected by non-ATS first parties. Thus, it seems this service is not significantly altering their behaviors by age or while the user is logged out.

Roblox’s privacy policy at the time these data were collected does not fully explain what we have observed. The child and adolescent data types shared with third party ATS are particularly concerning. Roblox stated the following in their privacy policy: “We may share non-identifying data of all users regardless of their age” for purposes such as marketing, reporting requirements, and service analytics [96]. However, for the child and adolescent traces, we observed data types such as personal identifiers being shared with third party ATS, and we are skeptical that these data shared are “non-identifying” and are being used for the stated purposes considering the destinations include known third party ATS. Additionally,

related to CCPA regulations, they stated that they “have no actual knowledge of selling or sharing the Personal Information of minors under 16 years of age” [96]. It is also important to note that at the time of data collection, we did not encounter any parental consent process when creating the child accounts for this service, which means parental consent was not obtained for the data processing we observed. Considering what we have observed, we remain skeptical of these behaviors and policy disclosures.

TikTok: For the child and adolescent categories, we observed that all data type categories were collected by first parties, both non-ATS and ATS. Also, for both the child and adolescent traces, data flows including device identifiers and user communications were shared with both non-ATS and ATS third parties. The adolescent category also included data flows to third party ATS containing user interests and behaviors data. Comparing across age categories, the adult category contains more data flows to third parties, but overall the three age groups’ observed data flows are very similar.

Considering the data types shared with third party ATS for the child and adolescent users, including device identifiers, user communications, and user interests and behaviors, TikTok’s privacy policy stated the following: “We may share the information that we collect with our corporate group or service providers as necessary for them to support the internal operations of the TikTok service” [107]. However, we are skeptical that “service providers” would include third party ATS, which are known to be related to non-functional advertising and tracking purposes. Additionally, the policy stated, “TikTok does not sell information from children to third parties and does not share such information with third parties for the purposes of cross-context behavioral advertising” [107]. These statements lead us to question whether these observed practices of sharing data with third party ATS are for cross-context behavioral advertising or monetary consideration, since we observed both identifiers and personal information being shared with third party ATS. Also, it is important to note that at the time of data collection, we did not encounter any parental consent process when creating the child accounts for this service, and thus parental consent was not obtained for the data processing we observed.

YouTube: Both the regular YouTube service (i.e., for adolescents and adults) and YouTube Kids are included in the YouTube category, and the child trace in Table 4 corresponds to YouTube Kids specifically. We observed that all listed data type categories were collected by non-ATS first parties, and we observed no data flows to any third parties. Since Google owns YouTube and many ATS domains, it is understandable that we did not see any other third party domains contacted. The data flows in the table for the adolescent and adult groups were the same, with all listed data type categories collected by first party ATS. For the child trace, we observed that all listed data types except for personal identifiers and geolocation were collected by first party ATS.

In contrast with the other services discussed, YouTube’s privacy policies seem to be consistent with the behaviors we observed. For example, the YouTube Kids privacy policy stated that they collect information including device type and settings, log information, and unique identifiers from users for various reasons, such as “internal operational purposes”, “personalized content”, and “contextual advertising, including ad frequency capping” [33]. It appears the

observed behaviors are in line with the data collection and purposes disclosed for young users.

Platform Differences: Table 4 also shows whether each data flow was observed on the mobile app, website, neither, or both platforms for each service. Since we were able to collect more clear text traffic on the websites compared to the mobile apps, it is understandable that we see far more website-only data flows than mobile-only, but we do see many data flows that were observed on both platforms. We observed data flows unique to the mobile platform only for Roblox, TikTok, Minecraft, and Duolingo, and these mobile-only data flows largely involved sharing data with third parties, including some third party ATS, such as personal and device identifiers, personal characteristics, geolocation, user communications, and user interests and behaviors. We observed many data flows unique to the website platform for all services and across all data types. These results suggest some potential development differences in mobile apps and their website counterparts that could lead to differences in what data is collected/shared about users on each platform.

Key Takeaways: All but one of the services engaged in behaviors that raise concern for young users' privacy, including collecting and sharing both identifiers and personal information with third party ATS domains, which were often unclear and/or undisclosed in the services' privacy policies, while logged in and out. No service exhibited significantly different data processing treatment of the child and adolescent users compared to the adult users.

4.2 Data Linkability

We also investigate data linkability in the observed data flows per age group and while logged out. Data linkability refers to data types that are sent to the same third parties, including both ATS and non-ATS, that could be linked and used for tracking and profiling of users, as discussed in an SoK by Powar et al. [54]. In our case, data linkability could occur if data flows containing at least one data type from both the identifiers and personal information categories are sent to the same third party (e.g., a device hardware identifier and app/service usage data).

Across the entire dataset (not just linkable data flows), we found that the observed domains correspond to at least 212 companies (for some domains, we could not determine the owner). Based on our categorization of first/third parties and ATS domains, the destinations contacted include 320 first parties (e.g., tiktok.com, roblox.com, and duolingo.com), 33 first party ATS (e.g., clarity.ms, metrics.roblox.com, and browser.events.data.microsoft.com), 150 third parties (e.g., cloudfront.net, googleapis.com, and vimeocdn.com), and 485 third party ATS (e.g., google-analytics.com, amazon-adssystem.com, and doubleclick.net).

Figure 3 shows the counts of third parties that were specifically sent linkable data types based on the observed network traffic per service and per trace category. We observe that Quizlet had the highest counts of third parties to whom they sent linkable data for all categories except for the child trace, and YouTube has the least (all counts are zero), as YouTube only contacted first parties in our traces. We also observed a trend in the counts of third parties when comparing the age groups, with most of the services sharing linkable data types with a smaller number of third parties for

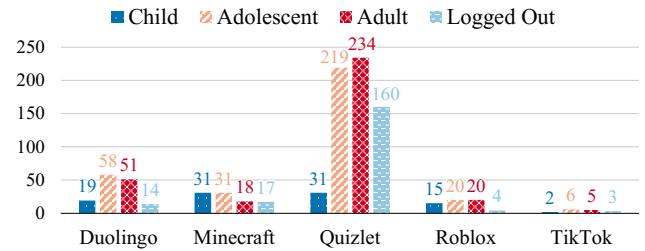


Figure 3: Counts of Third Parties Sent Linkable Data Types Per Service and Trace Category. Counts include third party domains, both ATS and non-ATS, that were sent linkable data types from each service per trace category (*i.e.*, child, adolescent, adult, and logged out).

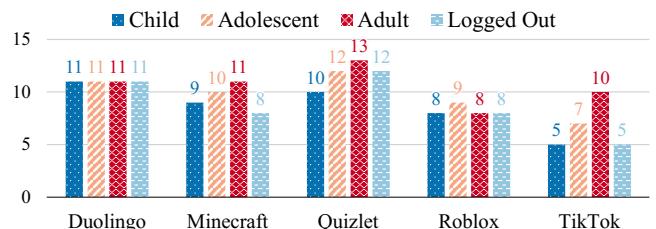


Figure 4: Sizes of Largest Sets of Linkable Data Types. A set of linkable data types is defined as all the data types that were shared with third party domains, including both ATS and non-ATS. The graph shows the size of the largest set of linkable data types shared by each service per trace category (*i.e.*, child, adolescent, adult, and logged out).

the child category compared to that of the adolescent and adult categories. However, we observed high counts for the adolescent category similar to those of the adult across services (*e.g.*, 219 for adolescent vs. 234 for adult in Quizlet).

We also analyze the sets of linkable data types transmitted by the services, and we show the sizes of the largest sets per trace and service in Figure 4. The largest set of linkable data types out of the entire dataset was shared by Quizlet in the adult trace, which included 13 (out of 19 total observed in the entire dataset) data types: network connection information, language, device information, app or service usage, service information, products and advertising, account settings, aliases, name, login information, location time, device software identifiers, and reasonably linkable personal identifiers. The most common set of linkable data types shared across the entire dataset included 5 data types: network connection information, language, device information, app or service usage, and service information.

To give context for the particular ATS with whom these services shared linkable data per trace category, Figure 5 shows an alluvial diagram of the most contacted third party ATS domains per service, mapped with the organizations that own those domains, with whom linkable data was shared. Due to space, we only show the top 10 most contacted ATS per service and per category, since there are over 190 third party ATS domains. The bar widths indicate the

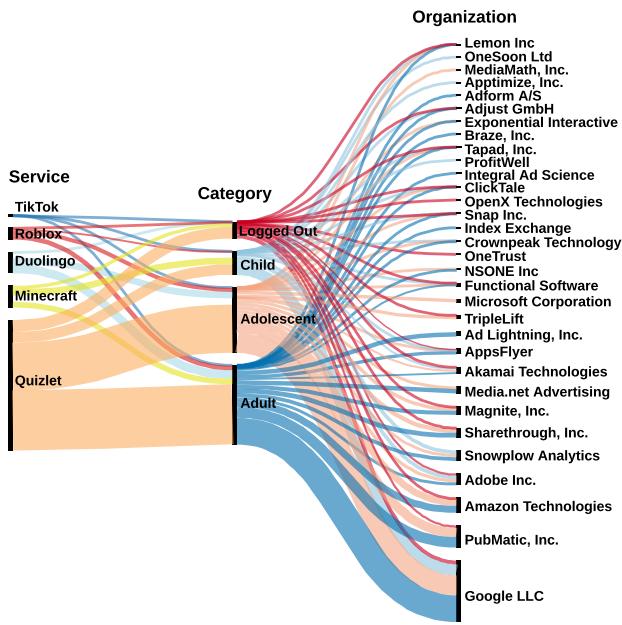


Figure 5: Most Frequent Third Party ATS Domains Sent Linkable Data Types. Alluvial diagram visualizes the top-10 most contacted third party ATS, shown based on their organizations, that were sent linkable data types by each service per trace category (*i.e.*, child, adolescent, adult, and logged out).

frequency of contact with those domains, demonstrating the relative difference in frequency of linkable data flows between services, with Quizlet contacting the most ATS domains with linkable data flows most frequently, and TikTok contacting the least for both. The 32 organizations shown include Google, Pubmatic, Amazon, Adobe, and Microsoft.

Key Takeaways: All services except one sent linkable data types to third party domains, including ATS and non-ATS, for all age groups and while logged out. These services sent multiple linkable data types to many of the same third party domains, demonstrating a lack of differentiation between age groups and while logged out, as well as risky linkable data sharing behaviors overall.

5 Discussion

5.1 Privacy Implications

All the services we investigated in this work with DIFFAUDIT collected and/or shared personal information about child and adolescent users in such a way that raises concerns with respect to COPPA and CCPA. When comparing between the child, adolescent, and adult trace data flows with our differential analysis approach, we found that these services do not appear to be altering their data collection and sharing practices as much as is expected for child and adolescent users. Every service also collected and/or shared information about users prior to knowing their age and obtaining consent (*i.e.*, while logged out), including with third party ATS, which is inappropriate under COPPA and CCPA regulations. General audience services should wait until they are informed of user

age and receive the appropriate consent so that they proceed with the appropriate data collection and sharing for the age category. Also, all but one of the services shared data with third parties that could enable data linkage of children and adolescent users. Finally, all but one of the services had privacy policies that were inconsistent with the data flows we observed for the child and adolescent traces.

Our findings complement the work of Reyes et al. [47], which is the closest work to our study. Reyes et al. investigated COPPA compliance of Android apps for children and found that a majority of children's apps studied had potential COPPA compliance issues largely because of third-party SDKs used in the apps. Also, 19% of the apps they studied collected identifiers and personal information through SDKs. However, as we discussed in Section 3.2.2, our work identifies considerably more data types (up to 35 based on our ontology), whereas [47] focused on a smaller set of data types (*e.g.*, device identifiers, contact information, and geolocation). Additionally, we are able to explore a richer set of data collection and sharing practices because we study both mobile apps and websites in context of COPPA and CCPA, further advancing our knowledge about how online services handle children's and adolescents data.

5.2 Recommendations

Based on our observations, children's and adolescents' online privacy are not being adequately protected. General audience services need to improve their practices to protect their young users and ensure they are compliant with COPPA and CCPA. We urge service providers to scrutinize every piece of data being collected and shared before and after users log in to their service. Service providers should also be more transparent in their privacy policies and correctly disclose their data collection and sharing behaviors to ensure users and parents are able to provide informed consent.

5.3 Limitations & Future Directions

Labeling. One bottleneck that exists in machine learning-based approaches for data type labeling is the size and quality of the ground truth dataset. In this work, we relied on human experts to manually label a sample of our dataset to validate our final models. Future directions can include automated approaches for labeling ground truth datasets [42], crowdsourcing labeling, and using additional datasets as they become available. Additionally, our GPT-4-based labeling method can be used to train smaller models using the labeled data produced by our GPT-4 data type classification (*i.e.*, teacher-student training), and these models can be applied without the use of APIs, which can be costly depending on the scale. We will continue to evolve our data type classification as large language models continue to improve.

Interaction with Apps and Websites. Automating the interaction with and exploration of apps and websites is a problem on its own and presents a bottleneck in this type of research. Prior works have explored different methods for automating this interaction on different platforms [43, 102, 121]. In this study, we chose to manually interact with the services to increase the quality and depth of our dataset. However, even if the interaction with the services becomes automated for scalability reasons, the DIFFAUDIT framework can easily be applied. Additionally, due to the limited

exploration time with each service, there may be more features and options that become available over a longer period of usage, and thus there may be more data flows not observed in our dataset. Such limitations in exploration can be an issue with both manual and automated interaction methods, though we attempted to mitigate this by exhaustively exploring all available features during data collection.

Encrypted Data. When decryption of the network traffic was not possible, we still included encrypted data in our analyses of the destinations contacted by the services. There is likely more data types being collected and shared than we discussed that we could not observe in the encrypted data.

Privacy Laws. This study is focused on US privacy laws (COPPA and CCPA). However, the DIFFAUDIT framework can be easily expanded to scale in broader contexts by including other privacy laws in other jurisdictions, such as the GDPR [28]. This can be accomplished by incorporating other privacy laws' data type definitions in our data type ontology and conducting compliance analyses based on their specific regulations. Expanding DIFFAUDIT in this way can allow auditors to analyze the same services across various jurisdictions.

Other Applications. At the time of this study, there were only six general audience services among the top-100 most popular that fit our criteria: Duolingo, Minecraft, Quizlet, Roblox, TikTok, and YouTube/YouTube Kids. However, our DIFFAUDIT framework can be applied to other general audience services, as they become available in the future, since DIFFAUDIT relies only on the availability of network traffic. Furthermore, since DIFFAUDIT is platform-agnostic and relies only on network traffic, it can be applied to audit data collection and sharing practices of online services for young users on emerging platforms, such as virtual reality applications, voice assistants, and smart toys. Our differential analysis approach can also be re-purposed for other age-based auditing, such as content analysis and dark patterns.

Using DIFFAUDIT. We envision that DIFFAUDIT can be used by researchers and regulators to identify potentially problematic behaviors and to audit compliance of general audience online services with COPPA and CCPA.

5.4 Ethics Considerations

We created test accounts on online services and collected network traffic while the authors interacted with those services. This may lead to advertisers losing some ad revenue, since ads may show up as we use the services in our data collection sessions, but we only use each service as long as is required to exhaust all the functionalities, which is not long, thus minimizing our impact. However, no real users' data were collected or used. We plan to reach out to the online service providers and notify them of our observations and findings (responsible disclosure).

6 Conclusion

In this work, we present DIFFAUDIT, a platform-agnostic auditing framework for general audience services to investigate privacy behaviors and regulatory compliance. Within this framework, we construct data flows from network traffic through destination analysis and our novel data type classification method with GPT-4 and

our data type ontology, which identifies considerably more data types than prior work. We conduct differential analyses to investigate the observed data flows generated by child, adolescent, and adult user profiles. We compare the data flows against regulations and privacy policies, analyze the data flows while logged in and out (*i.e.*, before and after consent is given and user age is disclosed), and investigate data linkability. We apply DIFFAUDIT to popular online services on mobile and website platforms and find that all audited services engage in problematic behaviors for children and adolescents, including collecting and/or sharing data while users are logged out, inconsistent privacy policy disclosures, and data flows shared with third parties that enable data linkage for children and adolescents. DIFFAUDIT enables auditors to investigate network traffic through differential analysis, utilizing an automated and customizable data type classification method, to identify suspicious behaviors in network traffic for any platform with a focus on children's and adolescents' privacy under COPPA and CCPA.

We make DIFFAUDIT's implementation and datasets available at <https://athinagroup.eng.uci.edu/projects/childrens-privacy/>.

Acknowledgements

We thank our shepherd and reviewers for their helpful feedback. This work is supported both by NSF awards FG22924 and FG22490 and a gift from the UC Noyce Initiative.

References

- [1] Abbas Razaghpanah, Rishab Nithyanand, Narseo Vallina-Rodriguez, Srikanth Sundaresan, Mark Allman, Christian Kreibich, and Phillipa Gill. 2018. Apps, Trackers, Privacy, and Regulators: A Global Study of the Mobile Tracking Ecosystem. In *Proceedings 2018 Network and Distributed System Security Symposium*, Vol. NDSS2018. Internet Society, San Diego, CA, 15. <https://doi.org/10.14722/ndss.2018.23353>
- [2] Adam D. Thierer. 2007. Social Networking and Age Verification: Many Hard Questions; No Easy Solutions. *SSRN Electronic Journal* 14, 5 (2007), 33. <https://doi.org/10.2139/ssrn.976936>
- [3] Alem Fitwi, Meng Yuan, Seyed Nikouei, and Yu Chen. 2020. Minor Privacy Protection by Real-time Children Identification and Face Scrambling at the Edge. *ICST Transactions on Security and Safety* 7, 23 (June 2020), 164560. <https://doi.org/10.4108/eai.13-7-2018.164560>
- [4] Anastasia Shuba, Anh Le, Emmanouil Alimpertis, Minas Gjoka, and Athina Markopoulou. 2017. AntMonitor: A System for On-Device Mobile Network Monitoring and its Applications. <http://arxiv.org/abs/1611.04268> arXiv:1611.04268 [cs].
- [5] Anastasia Shuba and Athina Markopoulou. 2020. NoMoATS: Towards Automatic Detection of Mobile Tracking. *Proceedings on Privacy Enhancing Technologies* 2020, 2 (April 2020), 45–66. <https://doi.org/10.2478/popets-2020-0017>
- [6] Anastasia Shuba, Athina Markopoulou, and Zubair Shafiq. 2018. NoMoAds: Effective and Efficient Cross-App Mobile Ad-Blocking. *Proceedings on Privacy Enhancing Technologies* 2018, 4 (Oct. 2018), 125–140. <https://doi.org/10.1515/popets-2018-0035>
- [7] Anirudh Ekambaranathan, Jun Zhao, and Max Van Kleek. 2021. “Money makes the world go around”: Identifying Barriers to Better Privacy in Children’s Apps From Developers’ Perspectives. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–15. <https://doi.org/10.1145/3411764.3445599>
- [8] Arun Kumar Ghosh, Karla Badillo-Urquiola, Shion Guha, Joseph J. LaViola Jr, and Pamela J. Wisniewski. 2018. Safety vs. Surveillance: What Children Have to Say about Mobile Apps for Parental Control. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, Montreal QC Canada, 1–14. <https://doi.org/10.1145/3173574.3173698>
- [9] Benjamin Andow, Samin Yaseer Mahmud, Justin Whitaker, William Enck, Bradley Reaves, Kapil Singh, and Serge Egelman. 2020. Actions Speak Louder than Words: Entity-Sensitive Privacy Policy and Data Flow Analysis with PoliCheck. In *29th USENIX Security Symposium*. USENIX Association, Virtual, 985–1002. <https://www.usenix.org/conference/usenixsecurity20/presentation/andow>
- [10] Bing Hu, Bin Liu, Neil Zhenqiang Gong, Deguang Kong, and Hongxia Jin. 2015. 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, Melbourne Australia, 1111–1120. <https://doi.org/10.1145/2806416.2806579>

[11] Brooke Auxier, Monica Anderson, Andrew Perrin, and Erica Turner. 2020. *Parenting Children in the Age of Screens*. Technical Report. Pew Research Center. <https://www.pewresearch.org/internet/2020/07/28/childrens-engagement-with-digital-devices-screen-time/>

[12] California Legislative Information. 2018. California Consumer Privacy Act of 2018. 1798 C.A. Civil § 120 (2018). https://leginfo.legislature.ca.gov/faces/codes_displayText.xhtml?division=3.&part=4.&lawCode=CIV&title=1.81.5

[13] California Legislative Information. 2018. California Consumer Privacy Act of 2018. 1798 C.A. Civil § 140 (2018). https://leginfo.legislature.ca.gov/faces/codes_displayText.xhtml?division=3.&part=4.&lawCode=CIV&title=1.81.5

[14] Cami Goray and Sarita Schoenebeck. 2022. Youths' Perceptions of Data Collection in Online Advertising and Social Media. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (Nov. 2022), 1–27. <https://doi.org/10.1145/3555576>

[15] Chelsea Jarvie and Karen Renaud. 2021. Are you over 18? A snapshot of current age verification mechanisms. In *2021 IFIP 8.11/11.13 Dewald Roode Information Security Research Workshop*. 2021 Dewald Roode Workshop, Antonio, Texas, USA, 25. <https://pureportal.strath.ac.uk/en/publications/are-you-over-18-a-snapshot-of-current-age-verification-mechanisms>

[16] Code of Federal Regulations. 2013. Children's Online Privacy Protection Rule. 16 C.F.R. § 312.2 (2013). <https://www.ecfr.gov/current/title-16/chapter-I/subchapter-C/part-312>

[17] Cristiana S. Silva, Glávia A.R. Barbosa, Ismael S. Silva, Tatiane S. Silva, Fernando Mourão, and Flávio Coutinho. 2017. Privacy for Children and Teenagers on Social Networks from a Usability Perspective: A Case Study on Facebook. In *Proceedings of the 2017 ACM on Web Science Conference*. ACM, Troy New York USA, 63–71. <https://doi.org/10.1145/3091478.3091479>

[18] Danah Boyd, Eszter Hargittai, Jason Schultz, and John Palfrey. 2011. Why parents help their children lie to Facebook about age: Unintended consequences of the Children's Online Privacy Protection Act*. *First Monday* 16, 11 (Oct. 2011), 27. <https://doi.org/10.5210/fm.v16i11.3850>

[19] DuckDuckGo. 2023. DuckDuckGo Tracker Radar. <https://github.com/duckduckgo/tracker-radar>

[20] Duolingo. 2023. Duolingo. <https://www.duolingo.com>

[21] Duolingo. 2023. Duolingo Privacy Policy. <https://www.duolingo.com/privacy>

[22] Anirudh Ekambaranathan, Jun Zhao, and Max Van Kleek. 2023. How Can We Design Privacy-Friendly Apps for Children? Using a Research through Design Process to Understand Developers' Needs and Challenges. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW2 (Sept. 2023), 1–29. <https://doi.org/10.1145/3610066>

[23] Elmira Deldari, Diana Freed, Julio Poveda, and Yaxing Yao. 2023. An Investigation of Teenager Experiences in Social Virtual Reality from Teenagers', Parents', and Bystanders' Perspectives. In *Nineteenth Symposium on Usable Privacy and Security (SOUPS 2023)*. USENIX Association, Anaheim, CA, 1–17. <https://www.usenix.org/conference/soups2023/presentation/deldari>

[24] Emanuele Faranda. 2023. PCAPdroid. <https://github.com/emanuele-f/PCAPdroid>

[25] Emanuele Faranda. 2023. PCAPdroid User Guide. https://emanuele-f.github.io/PCAPdroid/tls_decription

[26] Emily A. Vogels, Risa Gelles-Watnick, and Navid Massarat. 2022. *Teens, Social Media and Technology 2022*. Technical Report. Pew Research Center. <https://www.pewresearch.org/internet/2022/08/10/teens-social-media-and-technology-2022/>

[27] Emily McReynolds, Sarah Hubbard, Timothy Lau, Aditya Saraf, Maya Cakmak, and Franziska Roessner. 2017. Toys that Listen: A Study of Parents, Children, and Internet-Connected Toys. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, Denver Colorado USA, 5197–5207. <https://doi.org/10.1145/3025453.3025735>

[28] European Commission. 2024. Data protection in the EU. https://commission.europa.eu/law/law-topic/data-protection/data-protection-eu_en

[29] Federal Trade Commission. 2020. Complying with COPPA: Frequently Asked Questions. <https://www.ftc.gov/business-guidance/resources/complying-coppa-frequently-asked-questions>

[30] France Bélanger, Robert E. Crossler, Janine S. Hiller, Jung-Min Park, and Michael S. Hsiao. 2013. POCKET: A tool for protecting children's privacy online. *Decision Support Systems* 54, 2 (Jan. 2013), 1161–1173. <https://doi.org/10.1016/j.dss.2012.11.010>

[31] Ge Wang, Jun Zhao, Max Van Kleek, and Nigel Shadbolt. 2022. Informing Age-Appropriate AI: Examining Principles and Practices of AI for Children. In *CHI Conference on Human Factors in Computing Systems*. ACM, New Orleans LA USA, 1–29. <https://doi.org/10.1145/3491102.3502057>

[32] Ge Wang, Jun Zhao, Max Van Kleek, and Nigel Shadbolt. 2023. 12 Ways to Empower: Designing for Children's Digital Autonomy. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM, Hamburg Germany, 1–27. <https://doi.org/10.1145/3544548.3580935>

[33] Google. 2019. YouTube Kids Privacy Notice. <https://kids.youtube.com/t/privacynote>

[34] Google. 2023. Google Chrome. <https://www.google.com/chrome/>

[35] Google. 2023. Google Play Store. <https://play.google.com/store/games>

[36] Google. 2023. YouTube. <https://www.youtube.com>

[37] Google. 2023. YouTube Kids. <https://www.youtubekids.com>

[38] Gordon Chu, Noah Aphorpe, and Nick Feamster. 2019. Security and Privacy Analyses of Internet of Things Children's Toys. *IEEE Internet of Things Journal* 6, 1 (Feb. 2019), 978–985. <https://doi.org/10.1109/JIOT.2018.2866423>

[39] Hao Cui, Rahmadi Trimananda, Athina Markopoulou, and Scott Jordan. 2023. PoliGraph: Automated Privacy Policy Analysis using Knowledge Graphs. In *32nd USENIX Security Symposium (USENIX Security 23)*. USENIX Association, Anaheim, CA, 1037–1054. <https://www.usenix.org/conference/usenixsecurity23/presentation/cui>

[40] Haojin Jin, Minyi Liu, Kevan Dodhia, Yuanchun Li, Gaurav Srivastava, Matthew Fredrikson, Yuvraj Agarwal, and Jason I. Hong. 2018. Why Are They Collecting My Data?: Inferring the Purposes of Network Traffic in Mobile Apps. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 4 (Dec. 2018), 1–27. <https://doi.org/10.1145/3287051>

[41] Helen Nissenbaum. 2009. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford University Press, Stanford, CA. <http://www.sup.org/books/title/?id=8862>

[42] Hieu Le, Salma Elmalaiki, and Athina Markopoulou. 2023. AutoFR: Automated Filter Rule Generation for Adblocking. In *32nd USENIX Security Symposium (USENIX Security 23)*. USENIX Association, Anaheim, CA, 7535–7552. <https://www.usenix.org/conference/usenixsecurity23/presentation/le>

[43] Hooman Mohajeri Moghaddam, Gunes Acar, Ben Burgess, Arunesh Mathur, Danny Yuxing Huang, Nick Feamster, Edward W. Felten, Prateek Mittal, and Arvind Narayanan. 2019. Watching You Watch: The Tracking Ecosystem of Over-the-Top TV Streaming Devices. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. ACM, London United Kingdom, 131–147. <https://doi.org/10.1145/3319535.3354198>

[44] Hugging Face. 2023. facebook/bart-large-mnli. <https://huggingface.co/facebook/bart-large-mnli>

[45] Hugging Face. 2023. Zero-Shot Classification. <https://huggingface.co/tasks/zero-shot-classification>

[46] Hyunjin Kang, Wonsun Shin, and Junru Huang. 2022. Teens' privacy management on video-sharing social media: the roles of perceived privacy risk and parental mediation. *Internet Research* 32, 1 (Jan. 2022), 312–334. <https://doi.org/10.1108/INTR-01-2021-0005>

[47] Irwin Reyes, Primal Wijesekera, Joel Reardon, Amit Elazari Bar On, Abbas Razaghpanah, Narseo Vallina-Rodriguez, and Serge Egelman. 2018. "Won't Somebody Think of the Children?" Examining COPPA Compliance at Scale. *Proceedings on Privacy Enhancing Technologies* 2018, 3 (June 2018), 63–83. <https://doi.org/10.1515/popets-2018-0021>

[48] Mike Isaac. 2024. Six takeaways from a contentious online child safety hearing. <https://www.nytimes.com/2024/01/31/technology/tech-senate-hearing-child-safety.html>

[49] Jan Odvarko. 2007. HAR 1.2 Spec. <http://www.softwareishard.com/blog/har-12-spec/>

[50] Janus Varmarken, Hieu Le, Anastasia Shuba, Athina Markopoulou, and Zubair Shafiq. 2020. The TV is Smart and Full of Trackers: Measuring Smart TV Advertising and Tracking. *Proceedings on Privacy Enhancing Technologies* 2020, 2 (April 2020), 129–154. <https://doi.org/10.2478/popets-2020-0021>

[51] Jingjing Ren, Ashwin Rao, Martina Lindorfer, Arnaud Legout, and David Choffnes. 2016. ReCon: Revealing and Controlling PII Leaks in Mobile Network Traffic. In *Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services*. ACM, Singapore Singapore, 361–374. <https://doi.org/10.1145/2906388.2906392>

[52] John Dempsey, Gavin Sim, Brendan Cassidy, and Vinh-Thong Ta. 2022. Children designing privacy warnings: Informing a set of design guidelines. *International Journal of Child-Computer Interaction* 31 (March 2022), 100446. <https://doi.org/10.1016/j.ijcci.2021.100446>

[53] John Kurkowski. 2023. john-kurkowski / tldeextract. <https://github.com/john-kurkowski/tldeextract>

[54] Jovan Powar and Alastair R. Beresford. 2023. SoK: Managing risks of linkage attacks on data privacy. *Proceedings on Privacy Enhancing Technologies* 2023, 2 (April 2023), 97–116. <https://doi.org/10.56553/popets-2023-0043>

[55] Jun Zhao, Ge Wang, Carys Dally, Petr Slovák, Julian Edbrooke-Childs, Max Van Kleek, and Nigel Shadbolt. 2019. 'I make up a silly name': Understanding Children's Perception of Privacy Risks Online. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, Glasgow, Scotland, UK, 1–13. <https://doi.org/10.1145/3290605.3300336>

[56] Kaiwen Sun, Carlo Sugatan, Tanisha Afnan, Hayley Simon, Susan A. Gelman, Jenny Radesky, and Florian Schaub. 2021. "They See You're a Girl if You Pick a Pink Robot with a Skirt": A Qualitative Study of How Children Conceptualize Data Processing and Digital Privacy Risks. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–34.

<https://doi.org/10.1145/3411764.3445333>

[57] Kalhan Rosenblatt. 2024. Florida Gov. Ron DeSantis signs bill that bans children under 14 from having social media accounts. <https://www.nbcnews.com/tech/florida-ron-desantis-signs-bill-social-media-kids-ban-rcna144950>.

[58] Kayce Basques. 2019. Chrome DevTools Network: Inspect Network Activity. <https://developer.chrome.com/docs/devtools/network/>

[59] Konrad Kollnig, Pierre Dewitte, Max Van Kleek, Ge Wang, Daniel Omeiza, Helena Webb, and Nigel Shadbolt. 2021. A fait accompli? An empirical study into the absence of consent to Third-Party tracking in android apps. In *Seventeenth symposium on usable privacy and security (SOUPS 2021)*. USENIX Association, Virtual, 181–196. <https://www.usenix.org/conference/soups2021/presentation/kollnig>

[60] Pierre Laperdrix, Walter Rudametkin, and Benoit Baudry. 2016. Beauty and the Beast: Diverting Modern Web Browsers to Build Unique Browser Fingerprints. In *2016 IEEE Symposium on Security and Privacy (SP)*. IEEE, San Jose, CA, 878–894. <https://doi.org/10.1109/SP.2016.57>

[61] Victor Le Pochat, Tom Van Goethem, Samaneh Tajalizadehkhooob, Maciej Korczyński, and Wouter Joosen. 2023. Tranco. <https://tranco-list.eu/>

[62] Leah Zhang-Kennedy, Christine Mekhail, Yomna Abdelaziz, and Sonia Chiasson. 2016. From Nosy Little Brothers to Stranger-Danger: Children and Parents' Perception of Mobile Threats. In *Proceedings of the The 15th International Conference on Interaction Design and Children*. ACM, Manchester United Kingdom, 388–399. <https://doi.org/10.1145/2930674.2930716>

[63] Lewis Tunstall, Nils Reimers, Unso Eun Seo Jo, Luke Bates, Daniel Korat, Moshe Wasserblat, and Oren Pereg. 2022. Efficient Few-Shot Learning Without Prompts. <https://doi.org/10.48550/ARXIV.2209.11055> Publisher: arXiv Version Number: 1.

[64] Ilaria Liccardi, Monica Bulger, Hal Abelson, Daniel J. Weitzner, and Wendy Mackay. 2014. Can apps play by the COPPA Rules? In *2014 Twelfth Annual International Conference on Privacy, Security and Trust*. IEEE, Toronto, ON, Canada, 1–9. <https://doi.org/10.1109/PST.2014.6890917>

[65] Liliana Pasquale, Paola Zippo, Cliona Curley, Brian O'Neill, and Marina Mongiello. 2022. Digital Age of Consent and Age Verification: Can They Protect Children? *IEEE Software* 39, 3 (May 2022), 50–57. <https://ieeexplore.ieee.org/document/9295422/>

[66] Maarten Grootendorst. 2021. MaartenGr/PolyFuzz. <https://doi.org/10.5281/ZENODO.4461050>

[67] Mamta Akter, Amy J. Godfrey, Jess Kropczynski, Heather R. Lipford, and Pamela J. Wisniewski. 2022. From Parental Control to Joint Family Oversight: Can Parents and Teens Manage Mobile Online Safety and Privacy as Equals? *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW1 (March 2022), 1–28. <https://doi.org/10.1145/3512904>

[68] Maria Grazia Lo Cricchio, Benedetta E. Palladino, Androulla Eleftheriou, Annalaura Nocentini, and Ersilia Menesini. 2022. Parental Mediation Strategies and Their Role on Youths' Online Privacy Disclosure and Protection: A Systematic Review. *European Psychologist* 27, 2 (April 2022), 116–130. <https://doi.org/10.1027/1016-9040/a000450>

[69] Marisa Meyer, Victoria Adkins, Nalingna Yuan, Heidi M. Weeks, Yung-Ju Chang, and Jenny Radesky. 2019. Advertising in Young Children's Apps: A Content Analysis. *Journal of Developmental & Behavioral Pediatrics* 40, 1 (Jan. 2019), 32–39. <https://doi.org/10.1097/DBP.0000000000000622>

[70] Martin Hilbert, Drew P. Cingel, Jingwen Zhang, Samantha L. Vigil, Jane Shawcroft, Haoning Xue, Arti Thakur, and Zubair Shafiq. 2024. #BigTech @Minors: Social Media Algorithms Personalize Minors' Content After a Single Session, but Not for Their Protection. <https://doi.org/10.2139/ssrn.4674573>

[71] Max Van Kleek, Ilaria Liccardi, Reuben Binns, Jun Zhao, Daniel J. Weitzner, and Nigel Shadbolt. 2017. Better the Devil You Know: Exposing the Data Sharing Practices of Smartphone Apps. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, Denver Colorado USA, 5208–5220. <https://doi.org/10.1145/3025453.3025556>

[72] Microsoft. 2023. Microsoft U.S. State Data Privacy Laws Notice. <https://privacy.microsoft.com/en-US/ccpa>

[73] Microsoft. 2023. Minecraft. <https://www.minecraft.net/en-us>

[74] Mike Lewis, Yinhai Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. <https://doi.org/10.48550/ARXIV.1910.13461> Publisher: arXiv Version Number: 1.

[75] N. Nikiforakis, A. Kapravelos, W. Joosen, C. Kruegel, F. Piessens, and G. Vigna. 2013. Cookieless Monster: Exploring the Ecosystem of Web-Based Device Fingerprinting. In *2013 IEEE Symposium on Security and Privacy*. IEEE, Berkeley, CA, 541–555. <https://doi.org/10.1109/SP.2013.43>

[76] Nico Ebert, Tim Geppert, Joanna Strycharz, Melanie Knieps, Michael Höning, and Elke Brucker-Kley. 2023. Creative beyond TikToks: Investigating Adolescents' Social Privacy Management on TikTok. *Proceedings on Privacy Enhancing Technologies* 2023, 2 (April 2023), 221–235. <https://doi.org/10.56553/popeps-2023-0049>

[77] Noah Aphorpe, Sarah Varghese, and Nick Feamster. 2019. Evaluating the Contextual Integrity of Privacy Regulation: Parents' IoT Toy Privacy Norms Versus COPPA. In *28th USENIX Security Symposium (USENIX Security 19)*. USENIX Association, Santa Clara, CA, 123–140. <https://www.usenix.org/conference/usenixsecurity19/presentation/aphorpe>

[78] Ole André Vadla Ravnås. 2024. Frida. <https://frida.re/>

[79] OpenAI. 2023. Chat Completions API. <https://platform.openai.com/docs/guides/gpt/chat-completions-api>

[80] OpenAI. 2023. Create chat completion. <https://platform.openai.com/docs/api-reference/chat/create#temperature>

[81] OpenAI. 2023. GPT-4. <https://platform.openai.com/docs/models/gpt-4>

[82] OpenAI. 2023. GPT best practices. <https://platform.openai.com/docs/guides/gpt-best-practices>

[83] Pamela Wisniewski, Arup Kumar Ghosh, Heng Xu, Mary Beth Rosson, and John M. Carroll. 2017. Parental Control vs. Teen Self-Regulation: Is there a middle ground for mobile online safety?. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. ACM, Portland Oregon USA, 51–69. <https://doi.org/10.1145/2998181.2998352>

[84] Patrick C. K. Hung, Kamen Kanev, Farkhund Iqbal, David Mettrick, Laura Rafferty, Guan-Pu Pan, Shih-Chia Huang, and Benjamin C. M. Fung. 2017. A Study of Children Facial Recognition for Privacy in Smart TV. In *Computational Modeling of Objects Presented in Images. Fundamentals, Methods, and Applications*. Reneta P. Barneva, Valentin E. Brimkov, and João Manuel R.S. Tavares (Eds.). Vol. 10149. Springer International Publishing, Cham, 229–240. https://doi.org/10.1007/978-3-319-54609-4_17 Series Title: Lecture Notes in Computer Science.

[85] Peiyi Yang, Jie Fan, Zice Wei, Haoqian Li, Tu Le, and Yuan Tian. 2023. Towards Usable Parental Control for Voice Assistants. In *Proceedings of Cyber-Physical Systems and Internet of Things Week 2023*. ACM, San Antonio TX USA, 43–48. <https://doi.org/10.1145/3576914.3587491>

[86] Phoebe K. Chua and Melissa Mazmanian. 2021. What Are You Doing With Your Phone?: How Social Class Frames Parent-Teen Tensions around Teens' Smartphone Use. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–12. <https://doi.org/10.1145/3411764.3445275>

[87] Priya Kumar, Jessica Vitak, Marshini Chetty, Tamara L. Clegg, Jonathan Yang, Brenna McNally, and Elizabeth Bonsignore. 2018. Co-designing online privacy-related games and stories with children. In *Proceedings of the 17th ACM Conference on Interaction Design and Children*. ACM, Trondheim Norway, 67–79. <https://doi.org/10.1145/3202185.3202735>

[88] Priya Kumar, Shalmali Milind Naik, Utkarsha Ramesh Devkar, Marshini Chetty, Tamara L. Clegg, and Jessica Vitak. 2017. 'No Telling Passcodes Out Because They're Private': Understanding Children's Mental Models of Privacy and Security Online. *Proceedings of the ACM on Human-Computer Interaction* 1, CSCW (Dec. 2017), 1–21. <https://doi.org/10.1145/3134699>

[89] Proxyman. 2023. Proxyman. <https://proxyman.io>

[90] Quizlet. 2021. Quizlet Privacy Policy. <https://quizlet.com/privacy>

[91] Quizlet. 2023. Quizlet. <https://quizlet.com>

[92] Rahmadi Trimananda, Hieu Le, Hao Cui, Janice Tran Ho, Anastasia Shuba, and Athina Markopoulou. 2022. OVRseen: Auditing Network Traffic and Privacy Policies in Oculus VR. In *31st USENIX Security Symposium (USENIX Security 22)*. USENIX Association, Boston, MA, USA, 3789–3806. <https://www.usenix.org/conference/usenixsecurity22/presentation/trimananda>

[93] Reuben Binns, Ulrik Lyngs, Max Van Kleek, Jun Zhao, Timothy Libert, and Nigel Shadbolt. 2018. Third Party Tracking in the Mobile Ecosystem. In *Proceedings of the 10th ACM Conference on Web Science*. ACM, Amsterdam Netherlands, 23–31. <https://doi.org/10.1145/3201064.3201089>

[94] Richard Penman. 2023. whois. <https://github.com/richardpenman/whois>

[95] Roblox. 2023. Roblox. <https://www.roblox.com/>

[96] Roblox. 2023. Roblox Privacy and Cookie Policy. <https://en.help.roblox.com/hc/en-us/articles/115004630823-Roblox-Privacy-and-Cookie-Policy>

[97] Scott Babwah Brennen and Matt Perault. 2023. Keeping Kids Safe Online: How Should Policymakers Approach Age Verification? <https://www.thecgo.org/research/keeping-kids-safe-online-how-should-policymakers-approach-age-verification/>

[98] Siew Yong, Dale Lindskog, Ron Ruhl, and Pavol Zavarsky. 2011. Risk Mitigation Strategies for Mobile Wi-Fi Robot Toys from Online Pedophiles. In *2011 IEEE Third Int'l Conference on Privacy, Security, Risk and Trust and 2011 IEEE Third Int'l Conference on Social Computing*. IEEE, Boston, MA, USA, 1220–1223. <https://doi.org/10.1109/PASSAT/SocialCom.2011.194>

[99] Simone van der Hof and Sanne Ouburg. 2022. We Take Your Word for It' - A Review of Methods of Age Verification and Parental Consent in Digital Services. *European Data Protection Law Review* 8, 1 (2022), 61–72. <https://heinonline.org/HOL/P2h=hein.journals/edpl8&i=67>

[100] Natasha Singer. 2024. Bipartisan Bill Aims to Protect Children Online. <https://www.nytimes.com/2024/01/31/technology/congress-social-media-safety.html?searchResultPosition=4>

[101] State of California Department of Justice Office of the Attorney General. 2023. California Consumer Privacy Act (CCPA). <https://www.oag.ca.gov/privacy/ccpa>

[102] Steven Englehardt and Arvind Narayanan. 2016. Online Tracking: A 1-million-site Measurement and Analysis. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. ACM, Vienna Austria, 1388–1401. <https://doi.org/10.1145/2976749.2978313>

[103] Suzan Ali, Mounir Elgharabawy, Quentin Duchaussoy, Mohammad Mannan, and Amr Youssef. 2020. Betrayed by the Guardian: Security and Privacy Risks of Parental Control Solutions. In *Annual Computer Security Applications Conference*. ACM, Austin USA, 69–83. <https://doi.org/10.1145/3427228.3427287>

[104] The Associated Press. 2024. Georgia joins states seeking parental permission before children join social media. <https://www.nbcnews.com/news/us-news/georgia-joins-states-seeking-parental-permission-children-join-social-rena145589>

[105] The Firebog. 2023. The Big Blocklist Collection. <https://firebog.net>

[106] TikTok. 2023. TikTok. <https://www.tiktok.com/>

[107] TikTok. 2023. TikTok Children's Privacy Policy. <https://www.tiktok.com/legal/page/global/childrens-privacy-policy/en>

[108] Tinhinane Medjkounne, Oana Goga, and Juliette Senechal. 2023. Marketing to Children Through Online Targeted Advertising: Targeting Mechanisms and Legal Aspects. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*. ACM, Copenhagen Denmark, 180–194. <https://doi.org/10.1145/3576915.3623172>

[109] Trung Tin Nguyen, Michael Backes, Ninja Marnau, and Ben Stock. 2021. Share First, Ask Later (or Never)? Studying Violations of GDPR's Explicit Consent in Android Apps. In *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, Virtual, 3667–3684. <https://www.usenix.org/conference/usenixsecurity21/presentation/nguyen>

[110] Tu Le, Danny Yuxing Huang, Noah Apthorpe, and Yuan Tian. 2022. SkillBot: Identifying Risky Content for Children in Alexa Skills. *ACM Transactions on Internet Technology* 22, 3 (Aug. 2022), 1–31. <https://doi.org/10.1145/3539609>

[111] Uma Iqbal, Pouneh Nikkhah Bahrami, Rahmadi Trimananda, Hao Cui, Alexander Gamero-Garrido, Daniel J. Dubois, David Choffnes, Athina Markopoulou, Franziska Roesner, and Zubair Shafiq. 2023. Tracking, Profiling, and Ad Targeting in the Alexa Echo Smart Speaker Ecosystem. In *Proceedings of the 2023 ACM on Internet Measurement Conference*. ACM, Montreal QC Canada, 569–583. <https://doi.org/10.1145/3618257.3624803>

[112] U.S. Federal Trade Commission. 2023. Children's Online Privacy Protection Rule ("COPPA"). <https://www.ftc.gov/legal-library/browse/rules/childrens-online-privacy-protection-rule-coppa>

[113] Victor Le Pochat, Tom Van Goethem, Samaneh Tajalizadehkhoob, Maciej Koczcynski, and Wouter Joosen. 2019. Tranco: A Research-Oriented Top Sites Ranking Hardened Against Manipulation. In *Proceedings 2019 Network and Distributed System Security Symposium*. Internet Society, San Diego, CA, 15. <https://doi.org/10.14722/ndss.2019.23386>

[114] William Enck, Peter Gilbert, Seungyeop Han, Vasant Tendulkar, Byung-Gon Chun, Landor P. Cox, Jaeyeon Jung, Patrick McDaniel, and Anmol N. Sheth. 2014. TaintDroid: An Information-Flow Tracking System for Realtime Privacy Monitoring on Smartphones. *ACM Transactions on Computer Systems* 32, 2 (June 2014), 1–29. <https://doi.org/10.1145/2619091>

[115] Wireshark. 2023. Embedding decryption secrets in a pcapng file. <https://wiki.wireshark.org/TLS#embedding-decryption-secrets-in-a-pcapng-file>

[116] Wireshark. 2023. Wireshark. <https://www.wireshark.org>

[117] Xiaomei Cai. 2008. Advertisements and Privacy: Comparing For-Profit and Non-Profit Web Sites for Children. *Communication Research Reports* 25, 1 (Feb. 2008), 67–75. <https://doi.org/10.1080/08824090701831826>

[118] Xiaomei Cai and Xiaoquan Zhao. 2013. Online advertising on popular children's websites: Structural features and privacy issues. *Computers in Human Behavior* 29, 4 (July 2013), 1510–1518. <https://doi.org/10.1016/j.chb.2013.01.013>

[119] Yan Shvartzshnaider, Noah Apthorpe, Nick Feamster, and Helen Nissenbaum. 2019. Going against the (Appropriate) Flow: A Contextual Integrity Approach to Privacy Policy Analysis. *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing* 7 (Oct. 2019), 162–170. <https://doi.org/10.1609/hcomp.v7i1.5266>

[120] Yihang Song and Urs Hengartner. 2015. PrivacyGuard: A VPN-based Platform to Detect Information Leakage on Android Devices. In *Proceedings of the 5th Annual ACM CCS Workshop on Security and Privacy in Smartphones and Mobile Devices*. ACM, Denver Colorado USA, 15–26. <https://doi.org/10.1145/2808117.2808120>

[121] Yuanchun Li, Ziyue Yang, Yao Guo, and Xiangqun Chen. 2017. DroidBot: a lightweight UI-Guided test input generator for android. In *2017 IEEE/ACM 39th International Conference on Software Engineering Companion (ICSE-C)*. IEEE, Buenos Aires, 23–26. <https://doi.org/10.1109/ICSE-C.2017.8>

[122] Zengrui Liu, Umar Iqbal, and Nitesh Saxena. 2022. Opted Out, Yet Tracked: Are Regulations Enough to Protect Your Privacy? <https://doi.org/10.48550/ARXIV.2202.00885> Publisher: arXiv Version Number: 3.

[123] Álvaro Feal, Paolo Calciati, Narsee Vallina-Rodriguez, Carmela Troncoso, and Alessandra Gorla. 2020. Angel or Devil? A Privacy Study of Mobile Parental Control Apps. *Proceedings on Privacy Enhancing Technologies 2020*, 2 (April 2020), 314–335. <https://doi.org/10.2478/popets-2020-0029>

Appendices

A Data Type Ontology

In Section 3.2, we discuss how we label data types according to our data type ontology categories. In this appendix, we present the complete data type ontology in Table 5.

B Data Type Classification

In Section 3.2, we discuss data type classification methods, including using OpenAI's GPT-4 and experiments with alternative methods. In this appendix, we provide details about the prompt used with our GPT-4 classification method. Below is the prompt used for data type classification, and following the below text, we provide the category labels and data type examples from our ontology as part of the prompt:

"You are a text classifier for network traffic payload data. I am going to give you some categories and examples for each category. Then I will give you text sequences that I want you to categorize using the provided categories. The input texts were collected from network traffic payloads. Try to determine the meaning of the input texts and use the similarity of the categories and input texts to do the classification. For text with acronyms and abbreviations, use the meaning of the acronyms and abbreviations to do the classification. Provide an explanation for each classification in 15 words or less. Report a score of confidence on a scale of 0 to 1 for each categorization. Format your response exactly like this for each input text: <input text> // <category> // <score> // <explanation>."

Table 5: Data Type Ontology for Data Type Classification based on COPPA and CCPA. Level 3 is used for data type classification labels. Some labels have been grouped into level 4 due to space. See Table 2 for all 35 labels.

Level 1	Level 2	Level 3	Level 4
Identifiers [13, 16]	Personal Identifiers	Name	first and last name, user name
		Linked Personal Identifiers	social security number, driver's license number, state identification card number, passport number
		Contact Information	email address, telephone number, phone number
		Reasonably Linkable Personal Identifiers	IP address, unique pseudonym/alias
		Aliases	alias, online identifier, unique personal identifier, unique id, GUID (globally unique identifier), UUID (universally unique identifier)
		Customer Numbers	customer number, account name, insurance policy number, bank account number, credit card number, debit card number
		Login Information	password, login, authorization, authentication
	Device Identifiers	Device Hardware Identifiers	IMEI (international mobile equipment identity), MAC (media access control) address, unique device identifier, processor serial number, device serial number
		Device Software Identifiers	advertising identifier, cookie, pixel tag, beacon, tracking identifier
		Device Information	display, height, width, FPS (frames per second), browser, bitrate, ABR (adaptive bitrate), ABR bitrate map, speed, device, delay, OS (operating system), rate, screen, sound, memory, history, CPU (central processing unit), buffer, latency, download, load, frame, depth, download speed, render
Personal Information [13, 16]	Personal Characteristics	Protected Classifications	race, skin color, national origin, ancestry, age, language, religion, gender, sexual orientation, marital status, military or veteran status, medical conditions, genetic information, disabilities
		Biometric Information	DNA, images, voiceprint, patterns, rhythms, physical characteristics or descriptions
	Personal History	Personal History	employment, education, financial information, medical information
	Geolocation	Precise Geolocation	GPS (global positioning system) location, coordinates, postal address, latitude, longitude
		Coarse Geolocation	city, town, country, region
		Location Time	time, timestamp, timezone, time offset, date
	User Communications	Communications	audio, text, and video communications
		Contacts	people communication with, contact list
		Internet Activity	IP addresses communicated with, browsing history, search history
		Network Connection Information	request, response, DNS (domain name system), TCP (transmission control protocol), TLS (transport layer security), RTT (round trip time), TTFB (time to first byte), protocol, client, connection, key, payload, host, referer, telemetry, cache
	Sensors	Sensor Data	audio recordings, video recordings, other sensor data (e.g., thermal, olfactory)
	User Interests and Behavior	Products and Advertising	records of personal property, products or services, or those considered, consumer's interaction with an advertisement, advertisement/ad engagement, bid, analytics, marketing, third party, advertiser
		App or Service Usage	user interaction with an application, user interaction with a website, session, usage session, content, video, audio, video buffer, audio buffer, play, volume, avatar, behavior, action, event, data, status, duration, timing
		Account Settings	account, settings, consent, permission
		Service Information	server, SDK (software development kit), API (application programming interface), site, URL (uniform resource locator), domain, version, script, URI (uniform resource identifier), application, page, app, CDN (content delivery network), DOM (document object model)
		Inferences	user preferences, characteristics, psychological trends, predispositions, behavior, attitudes, intelligence, abilities, aptitudes, personality, purchase history, purchase tendency