

Targeted and Troublesome: Tracking and Advertising on Children’s Websites

WARNING: Contains potentially NSFW images

Zahra Moti
Radboud University

Asuman Senol
imec-COSIC, KU Leuven

Hamid Bostani
Radboud University

Frederik Zuiderveen Borgesius
Radboud University

Veelasha Moonsamy
Ruhr University Bochum

Arunesh Mathur
Independent Researcher

Gunes Acar
Radboud University

Abstract—On the modern web, trackers and advertisers frequently construct and monetize users’ detailed behavioral profiles without consent. Despite various studies on web tracking mechanisms and advertisements, there has been no rigorous study focusing on websites targeted at children. To address this gap, we present a measurement of tracking and (targeted) advertising on websites directed at children. Motivated by the lack of a comprehensive list of child-directed (i.e., targeted at children) websites, we first build a multilingual classifier based on web page titles and descriptions. Applying this classifier to over two million pages from the Common Crawl dataset, we compile a list of two thousand child-directed websites. Crawling these sites from five vantage points, we measure the prevalence of trackers, fingerprinting scripts, and advertisements. Our crawler detects ads displayed on child-directed websites and determines if ad targeting is enabled by scraping ad disclosure pages whenever available. Our results show that around 90% of child-directed websites embed one or more trackers, and about 27% contain targeted advertisements—a practice that should require verifiable parental consent. Next, we identify improper ads on child-directed websites by developing an ML pipeline that processes both images and text extracted from ads. The pipeline allows us to run semantic similarity queries for arbitrary search terms, revealing ads that promote services related to dating, weight loss, and mental health, as well as ads for sex toys and flirting chat services. Some of these ads feature repulsive, sexually-explicit and highly-inappropriate imagery. In summary, our findings indicate a trend of non-compliance with privacy regulations and troubling ad safety practices among many advertisers and child-directed websites. To ensure the protection of children and create a safer online environment, regulators and stakeholders must adopt and enforce more stringent measures. **Keywords** – online tracking, advertising, children, privacy

1. Introduction

The proliferation of online tracking for analytics, behavioral advertising, and marketing has resulted in over a decade’s worth of research into this ecosystem. Prior research has shown that not only is online tracking rampant on

the web [1] but that trackers use increasingly-invasive tracking mechanisms—e.g., third-party cookies, tracking pixels, evercookies, and browser fingerprinting [2], [3], [4], [1]—to relentlessly build detailed profiles of users across the web without any consent for targeted advertising.

Such privacy concerns aside, online advertising has shown to be problematic in other ways. Ads and ad networks are a vector for distributing ransomware, malicious programs, and cryptojackers—posing a serious security threat to users [5], [6], [7], [8], [9], [10], [11], [12], [13]. Ad networks also suffer from click fraud, which is estimated to reach \$100 billion in 2023 [14], [15]. Finally, online ads often contain clickbait, untrustworthy, or distasteful content that peddle software downloads, listicles, and health supplements—all of which users find problematic to their online experience [16].

While online tracking and targeted advertising pose a threat to users of all ages, children especially bear an acute cost. Children may not fully understand the consequences of online tracking and revealing their personal data online [17], [18], but they yield immense “pester power” to influence their parents’ purchase decisions [19]. Thus, children are an attractive target audience for advertisers and marketers alike [19], [20], as they are more vulnerable to persuasive advertising [21], [22], [23], and susceptible to harmful content [24], [25].

Despite the aforementioned evidence that suggests a differential impact on children, there is little empirical research on online tracking and advertising practices on children’s websites. The lack of a comprehensive and updated list of websites directed at children poses a major challenge for studying children’s websites. Previous large-scale internet measurement studies have relied on popular website lists such as Tranco [26] and Alexa [27] (before it shut down in 2021 [28]), but these lists may not specify website categories, and even when they do, the website categories may not be reliable and comprehensive [29], [30]. As a result, prior work [31], [23] has only examined online tracking on at most a hundred children’s websites and has been restricted in scope and methods—lacking a comprehensive investigation of both online tracking and advertising. To overcome this limitation, we built our own repository of

child-directed websites. We trained a text-based classifier that detects children’s websites using HTML metadata fields such as `<title>` and `<description>`. The classifier is based on a pre-trained multilingual model that we fine-tuned for our binary classification task. Applying the classifier to the Common Crawl dataset [32], we compiled a list of 2K manually verified child-directed websites.

To study several online tracking, ad targeting, and problematic ad practices, we crawl our list of 2K child-directed websites—varying the location (five vantage points) and form factors (desktop & mobile). Starting with ad targeting, we study the extent to which ads that appear on children’s websites are targeted—a practice that has come under increasing scrutiny both in the EU and the US [33], [34], [35]. We then present an exploratory analysis of ads from categories deemed problematic for children, such as dating, weight loss, mental health, and ads that contain racy content. Next, we turn to online tracking, which is a necessary ingredient of behavioral advertising. We study the ecosystem of trackers and specifically quantify the prevalence of trackers, cookies, and use of browser fingerprinting techniques such as Canvas, Canvas Font, AudioContext Fingerprinting, and WebRTC local IP discovery [1].

Our work is especially pertinent in light of impending regulatory changes. In the US, there have been calls [35] to update the Children’s Online Privacy Protection Act (COPPA) [36] in order to prohibit “internet companies from collecting personal information from users who are 13 to 16 years old without their consent” and to “ban targeted advertising to children and teens.” The current US President Joe Biden has called for a ban on collecting data on and serving targeted ads to children [34]; whereas in the EU, the upcoming Digital Services Act (DSA) will specifically prohibit ads targeted at children [33].

Our research seeks to offer empirical evidence on advertising and tracking practices employed on children’s websites by making the following contributions:

- Using a lightweight classifier based on web page metadata fields, we build a repository of child-directed websites and crawl them to measure tracking and advertising practices using multiple vantage points and form factors (desktop & mobile).
- We measure targeted ads using two ad vendors’ (Google and Criteo) ad disclosure statements, and find that targeting is enabled for 73% of the ads we could measure.
- Using text and images extracted from the ads, we detect *racy* ads, and ads about *weight loss*, *dating*, and *mental health* using semantic similarity search based on a lightweight, multilingual language model. While this content analysis is exploratory, our method enables human-in-the-loop investigations with arbitrary queries, and it paves the way for the automatic content analysis of ads.
- We also find ads linking to malicious content, improper ads of sex toys, dating services, and ads containing sexually suggestive images (Figure 1).

All the data and software from our study will be made available to researchers.¹

2. Related Work

2.1. Web tracking measurements

Over the past decade, several web privacy measurements have shown the scale and complexity of online tracking [37], [1], [38], [39], [40]. Research on *stateful* tracking has examined how unique tracking identifiers are stored on the client side [41] using cookies [42], [43], localStorage [2], cache (ETags) [2], or other client-side storage mechanisms. On the other hand, research on *stateless* tracking has examined the use of fingerprinting, a mechanism that exploits differences in browsers and devices to obtain a likely unique identifier [44]. Past research has shown that there are various fingerprinting vectors, including fonts, clock skew, GPUs, audio hardware, installed writing scripts and browser extensions, among others [45], [46], [47], [1], [48], [49], [50]. Research on defense against fingerprinting has contributed methods to detect fingerprinting, tracking and advertising [51], [4], [1], [38], [52], [53]. Our study borrows heuristics from prior work [1], [38] to detect fingerprinting scripts, and we use existing filter lists to identify trackers and advertisers.

2.2. Tracking & ads on children’s media

Motivated by the challenges posed by ads to children, Cai and Zhao [23] manually labeled ads displayed on 117 children’s websites. They found that 68% of the websites featured ads, and less than half complied with COPPA. The authors also argued that children are unlikely to distinguish many ads from the website’s original content. Vlajic et al. [31] investigated online tracking on twenty websites from Alexa’s “Kids and Teens” category [27] from two vantage points (EU & US). The authors manually analyzed the HTTP headers and quantified hidden images (i.e., likely tracking pixels) loaded from ads and analytics domains. Compared to this past work, we study orders of magnitude more websites, follow more rigorous tracking measurement methods, and compare results across different vantage points. Additionally, we automatically detect targeted ads using ad disclosure pages and present an exploratory analysis of the content of ads that appear on children’s websites.

Focusing on mobile platforms, Reyes et al. [54] dynamically analyzed around 6,000 free children’s Android apps and found that most apps violate COPPA due to their use of third-party SDKs.

2.3. Improper and malicious ads

A recent line of research has investigated the content of online ads. Zeng et al. [16] conducted a survey with 1,000

1. We share the list of identified child-directed websites and a sample of advertisement disclosures on <https://github.com/targeted-and-troublesome/>.

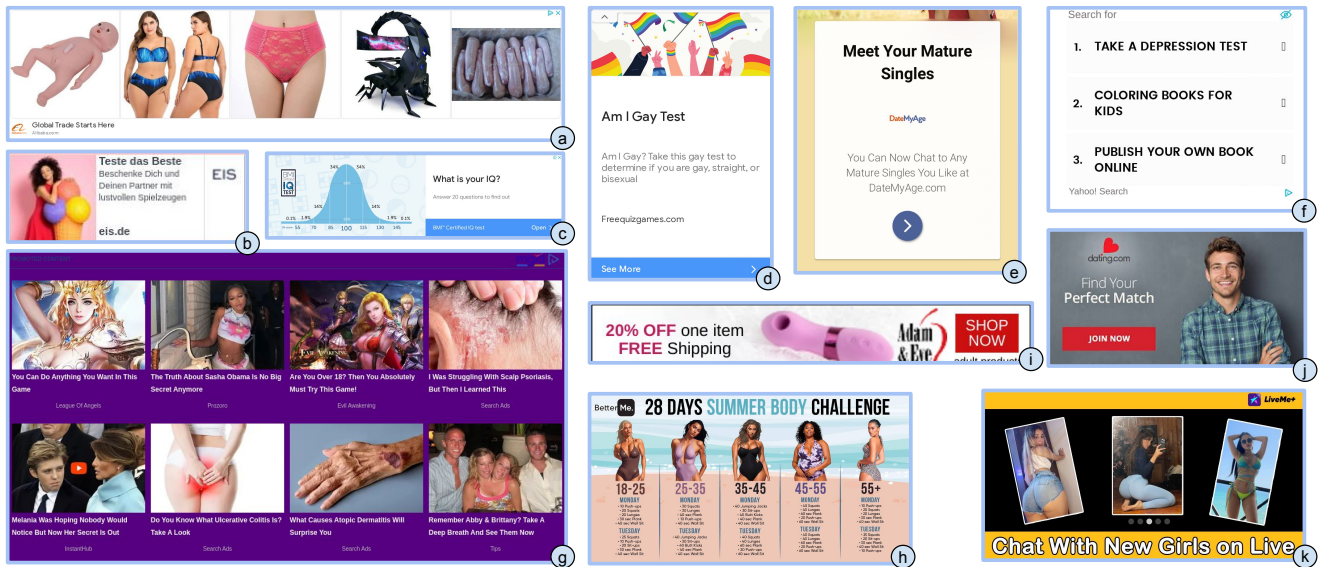


Figure 1: A sample of improper ads found on child-directed websites in our crawls.

participants to determine the type of advertising content (e.g., chumbox, clickbait, political, and low-quality content) that makes people dislike ads. In [55], the same authors also studied problematic ads in the news and misinformation websites, where they found problematic ads served by native ad platforms. In a study leading to the 2020 US elections, Zeng et al. [56] found that ads for misleading political polls that aim to hoover email addresses are widely used in online political advertising. Subramani et al. [7] studied the role of web push notifications in ad delivery, especially malicious ads. Through a large-scale study of malicious ads, Zarras et al. [5] showed that some ad exchanges are more prone to serving malicious ads due to inadequate detection. Akgul et al. [57] examined influencer VPN ads on YouTube and found advertisements disseminating misleading claims about online safety. Ali et al. [58] measured how the distribution of potentially harmful ads on Facebook varies across users. Venkatadri et al. [59] used Facebook’s advertiser interface to study how Facebook obtains personal identifiers used in advertising. In a concurrent work, Zhao et al. [60] analyzed mobile ads aimed at children, uncovering inappropriate advertisements by certified mobile ad SDKs. Medjkoune et al. [61] showed that advertisers can target ads to children by placing their ads in children-focused videos—bypassing YouTube’s age restrictions.

2.4. Ad transparency

In response to concerns about targeted advertising, ad networks and platforms have offered ad transparency interfaces that allow users to ascertain when and how they are being targeted. Andreou et al. [62] investigated Facebook’s ad explanations and found that they are often incomplete or misleading. Researchers have also argued that ad networks

should provide users with interpretable explanations and increase the visibility of disclosure mechanisms [63].

Bin Musa and Nithyanand [64] developed ATOM, a technique for determining data sharing between online trackers and advertisers. They used simulated personas to crawl websites, collect ad images, and conduct statistical analyses to identify correlations between tracker presence and advertiser behavior. Liu et al. [65] developed a framework called *AdReveal* to investigate different mechanisms used in online targeted advertising. Vallina et al. [66] used statements found in Google’s ad disclosure pages in their crowdsourced measurement of online behavioral advertisements. To detect stealthy ads that aim to bypass adblockers, Storey et al. [67] developed an extension that detects the AdChoices icon using perceptual hashing. While we considered applying Storey et al.’s method, we found URL-based detection of ad disclosure links (§4.6) to be more reliable and efficient.

2.5. Website categorization

The majority of studies on web categorization have focused on text-based classifiers because most web content and metadata are text-based [68], [69]. Various studies used machine learning models such as BERT and recurrent neural networks to learn contextual representations and features of web pages using meta tags and body content [70], [69], [68]. Other researchers proposed image-based web page classification techniques using pre-trained convolutional neural networks and using Google image search results [71], [72]. In our work, we built a lightweight classifier by fine-tuning an existing distilled language model and using text-based website metadata to detect child-directed websites.




Figure 2: Pipeline for building a list of child-directed websites.

3. Building a list of child-directed websites

It is estimated that there are more than one billion websites on the Internet [73], but only a small fraction are targeted at children. A central challenge, therefore, is identifying the websites that contain content directed to children. We initially searched and found three curated lists of children’s websites: kidSAFE Seal Program [74], CommonSense (filtered for children below the age of 13) [75], and a list compiled by Kaspersky [76]. Unfortunately, these lists contained only a total of 355 websites, some of which were no longer online.

To expand our limited list, we experimented with web categorization services such as McAfee, WebShrinker, and SimilarWeb, but decided to use VirusTotal’s (academic) API because other services were either not freely available or did not let us query in bulk. VirusTotal aggregates category labels from third-party scanners and categorization services, including BitDefender and TrendMicro [77]. We used the VirusTotal API to retrieve web category data for the top one million websites from the Chrome User Experience Report (CrUX) list from May 2022 [78]. We observed VirusTotal’s rate limits (20K/day/per academic license) during the process, which took roughly four weeks. By searching for substrings “kid” and “child” in the returned category labels and removing false positives (such as “Child abuse”), we obtained 1,264 websites categorized as related to children. However, our manual verification of these websites following the criteria presented in Appendix A revealed that 68.6% of them were false positives, yielding only 396 child-directed websites. Note that the low accuracy and inconsistency of domain classification/categorization services align with findings from prior work [30]. Combining our initial 355 websites with our verified list of 396 websites and removing all inaccessible (5) and duplicate (164) websites, we obtained a total of 582 child-directed websites.

Motivated by the lack of accurate, up-to-date, and comprehensive sources of child-directed websites, we built a classifier to detect child-directed websites using the list of 582 websites as labeled data. Figure 2 illustrates the training and fine-tuning process. We define “child-directed websites” as websites that are primarily intended for use by children, and contain content, activities, or other features that are likely to be appealing to children (see Appendix A, for our criteria). Note that our labeling criteria do not

fully overlap with COPPA’s definition [36]; for example, we do not require that sites have the actual knowledge of collecting children’s data. Thus, we do not claim to measure compliance with COPPA or other relevant laws. In our study, children refer to individuals under 13, aligning with US and EU regulations (§6.1).

3.1. Labeled data for ML classifier

Many web page classification methods use the entire text of the page [70] and its images [72], which can be resource-intensive and time-consuming. Alternatively, researchers have explored web page classification on metadata fields such as <title>, <description>, and <keywords>, which tend to be shorter and shown to have strong correlations with the topic of web pages [68]. We followed the latter approach for its computational efficiency and reasonable accuracy. Our preliminary analysis of over 500K web pages from the most popular one million websites in the Common Crawl dataset [32] showed that more than 97% of the websites have a title, 63% of the websites include a description, and 24% contain a keywords meta tag. Based on these availability statistics, we used the titles and descriptions for classification, leaving out the keywords. To extract the titles and descriptions, we used the following HTML tags: `title`, `description`, `og:[title|description]`, and `twitter:[title|description]`. Applying this method to the WAT metadata files from the June-July 2022 Common Crawl snapshot [32], we extracted the titles and descriptions, limiting ourselves to the top million websites in the Tranco [26] or the CrUX [78] list. We further refined our data by keeping a single page with the shortest URL from each hostname, which is more likely to be the home page. This resulted in metadata from 2.28 million pages, which included pages from the subdomains of the top million Tranco and CrUX websites. We also extracted the same title and metadata information from the 582 known child-directed websites using a simple script based on Playwright [79]. In both instances, when the page had more than one description or title available, we picked the longest one.

After completing the data collection process, we constructed a training set for our classifier. For negative samples, we randomly selected 2,500 of the 2.28 million pages and manually checked to remove children’s websites. Our positive samples consisted of 576 title-description pairs after filtering out websites with titles shorter than ten characters.

3.2. Building the ML classifier

Our training data contained a limited number of labeled samples and our input consisted of text-based meta fields, potentially in multiple languages. This made designing naive classifiers such as bag-of-words and TF-IDF less suitable for our task. Instead, we employed a pre-trained and multi-lingual language model. Pre-trained models have proven to be adequate for general text classification tasks, but

they need to be fine-tuned for more specific downstream tasks [70]. In particular, we decided to use the *Paraphrase-Multilingual-MPNet-base-v2* (PM-MPNet-v2) model from the *SentenceTransformers* [80], [81] library, which is a pre-trained multilingual and distilled model based on the MPNet method [82]. The distillation process [80], [83] involves training a smaller model (student) to mimic the behavior of a larger model (teacher). In particular, PM-MPNet-v2 is fine-tuned with a large number of paraphrased sentences in more than 50 languages [80]. However, PM-MPNet-v2 cannot be directly used for text classification since it only produces embeddings that are useful for semantic similarity-related tasks. Thus, we used HuggingFace’s *Trainer* API [84] and *AutoModelForSequenceClassification* [85] class to fine-tune the model and add a binary classification layer on top of the PM-MPNet-v2’s embedding outputs. As input to the classifier, we used the concatenation of title and descriptions since this combination gave the best accuracy compared to using title or description alone. In particular, we fine-tuned the model to detect child-directed websites using the training set explained in §3.1. We used the HuggingFace Transformers [86] and Ray Tune libraries’ Population Based Training (PBT) algorithm [87], [88] to find the best-performing hyperparameters (batch size=12, epochs=2, and learning rate=4.2e-05). The fine-tuning process took roughly five minutes on a consumer-grade GPU (GeForce RTX 3080 Ti).

To reduce false positives, we employed the modified “Classify-Verify” technique [89], which involves setting an acceptance threshold t , and accepting a prediction only if it is above t . Following Juarez et al. [90], we choose the threshold that maximizes $F_{\beta=0.5}$, which gives more weight to precision to reduce false positives. F_{β} is a weighted harmonic mean of precision and recall, adjustable to emphasize either metric according to the classification task’s needs [91]. A grid search of different threshold values shows that the maximum $F_{\beta=0.5}$ is achieved when $t = 0.93$, which reduces the false positive by 50%. Ultimately, our classifier achieved a precision of 86% and a recall of 70% using 10-fold cross-validation, as detailed in Table 1.

TABLE 1: Classification results before and after applying threshold (with 10-fold cross-validation).

	Precision	Recall	F-beta	TP	FP
Without threshold	0.79	0.81	0.79	47	12
With threshold	0.86	0.70	0.82	40	6

3.3. The list of 2K children’s websites

Using the fine-tuned classifier, we calculated the label and probability score for 2.28M web pages from Common Crawl, excluding websites used in the training process. This process took roughly 5 hours. Our classifier identified 53,092 web pages as children’s websites. Due to time constraints, we focused on manually verifying the top 2,500 websites sorted by classifier probability, starting with websites that are most likely to be child-directed. An evaluation

of our classifier and the details of our manual verification process can be found in Appendix A.1.

Our final list contained 2,004 websites in 48 distinct languages after eliminating false positives and deduplicating websites by their registrable domain (TLD+1). English was the most prevalent, accounting for 63% of all websites. The prevalence of other prominent languages, including Russian, Spanish, French, German, and Portuguese, ranging between 3% and 6%. The list included 582 websites from the training data and 1,422 websites identified by the classifier.

Website ranks: 1,422 of the 2,004 websites were ranked in the top 1 million Tranco list (median rank 304K). While over a quarter of the websites are in the top 200K ranks, websites from all popularity levels are captured in our list. 404 of the 582 websites that are not ranked by Tranco were ranked in the top one million by the CrUX list. Only 163 (8%) websites were not ranked either by Crux or Tranco in the top one million.

DNS0 Kids filter check: DNS0 Kids [92] is a domain name resolver that detects and filters out content that is not suitable for children such as adult, dating, and copyright-infringing content. To find out the status of the websites in our list, we compared DNS0 Kids with CloudFlare’s DNS resolver. If a website can be resolved by CloudFlare, but not by DNS0, we treated it as blocked. We found that only ten (0.5%) of the 2,004 websites in our list were blocked by DNS0. Reviewing these ten websites, we found six of them to contain pirated videos, including cartoons. The remaining four websites contained activities for children but it was not obvious to us why they were blocked.

4. Web Tracking and Advertising Measurements

To assess the prevalence of trackers, fingerprinting scripts, and (targeted) advertisements on child-directed websites, we extended Tracker Radar Collector (TRC) [93]. TRC is a Puppeteer-based [94] web crawler, which consists of modules called *collectors* that record different types of data during a crawl, such as HTTP requests/responses, cookies, screenshots, and JavaScript API calls. New collector modules can be easily added to collect data necessary to perform different web measurements such as ours. Specifically, we added the following collectors to TRC:

- `FingerprintCollector` (4.1): detects fingerprinting related function calls and property accesses
- `LinkCollector` (4.3): extracts inner page links
- `VideoCollector` (4.5): captures the crawl video
- `AdCollector` (4.6): detects ads and scrapes ad disclosures

We also used the existing TRC collectors, including `RequestCollector` to capture request/response details and detect tracking-related requests (4.2), `TargetCollector` to detect newly-opened tabs in 4.6, `CookieCollector` to analyze cookies, and finally `CMPCollector` (4.4) to interact with the consent dialogs and con-

sent management platforms (CMP). We used TRC’s anti-bot measures [93], which thwarts bot detection to a certain extent by overwriting artifacts typically probed by anti-bot scripts (e.g., `navigator.plugins`, `Notification.permission`) [95]. For the mobile crawls, we emulated a mobile browser using TRC’s built-in features to spoof viewport dimensions, touch support, and the user-agent string.

4.1. Identifying fingerprinting attempts

Identifying fingerprinting scripts can be challenging due to obfuscation and potential false positives. For example, scripts may use Canvas API for both drawing images or fingerprinting the user’s browser [47]. We draw on well-established methods to distinguish between fingerprinting and benign use of fingerprinting vectors [38], [1]. Specifically, we focused on Canvas, WebRTC, Canvas Font, or AudioContext fingerprinting and detected them using the heuristics presented by Iqbal et al. [38]. To detect fingerprinting attempts, we modified the `getter` and `setter` methods of the several Web APIs such as `CanvasRenderingContext2D.fillText` and `HTMLCanvasElement.toDataURL` to intercept potentially fingerprinting-related function calls and property accesses. Although TRC can intercept JavaScript API calls, we implemented a separate collector (`FingerprintCollector`) to avoid a known issue that prevented TRC from intercepting early function calls [96]. `FingerprintCollector` simply injects the instrumentation script into each page and its descendant frames as soon as they are created. We verified that our collector captures calls missed by TRC on both our custom-developed fingerprinting test pages and external demo pages such as `BrowserLeaks` [97].

4.2. Identifying tracking-related requests

To identify tracking-related requests, we utilized the `uBlock Origin Core` [98] npm package, which mimics the tracking protection of the widely-used `uBlock Origin` extension [99]. We used `uBlock Origin`’s default filter lists, which include `EasyList` and `EasyPrivacy`, among others [100]. To accurately detect tracking-related requests, we provided `uBlock Origin Core` with the request’s resource type (e.g., image or script) and the page and request URL derived from the HTTP request/response details recorded by the crawler. Then, we mapped the tracker domains to their owner entities (i.e., organizations/companies) using `DuckDuckGo`’s entity map [101]. Using entities to quantify tracker prevalence reduces overcounting as multiple domains can be owned by the same business (e.g., `googleanalytics.com` and `doubleclick.net` are both owned by Google).

4.3. Discovering inner pages

We refrained from only focusing on homepages as prior work found that websites’ inner pages tend to contain more

trackers and cookies [102], [103]. Thus, we also gathered five inner links from each of the 2,004 websites by conducting four separate link-collection crawls (desktop and mobile crawls from Frankfurt and NYC). We preferred to crawl sites from two vantage points to reduce the time and effort required for the link collection process. We excluded external domain links and documents such as PDFs and images, and we preferred links near the viewport’s center to avoid unrelated links from footers or less visible page areas. After gathering these inner links, we merged them with the homepage URLs to form the final set.

4.4. Interacting with consent dialogs

Since the GDPR came into effect, websites typically show consent dialogs when viewed from the EU and to some extent even from the US [104]. Ignoring these dialogs may lead to *undermeasurement* of the tracking and advertising practices. We decided to provide affirmative consent to all data processing request options (accept all) in our crawls to measure the full extent of advertisements and tracking a child could experience. To handle consent dialogs in an accurate and automated manner, we used `DuckDuckGo`’s `autoconsent` library [105], which comes bundled with TRC [106]. `Autoconsent` incorporates rules from `Consent-O-Matic` [107], [108], allowing programmatic interactions with the CMPs.

4.5. Video screen captures

To detect ads and scrape their disclosures, our crawler performed a series of interactions with the page, including dismissing popup dialogs, interacting with CMPs, and clicking on visible ad disclosure links (§ 4.6). To monitor these interactions, we added a video capture functionality to the crawler (`VideoCollector`). We used videos of the crawler’s interactions to troubleshoot potential issues with the crawler process as well as to label animated ads and other crawl artifacts manually.

4.6. Identifying ads and ad targeting criteria


The `AdCollector` performed three main functions: 1) detecting ads, 2) scraping ads—including its screenshots, links, iframes, scripts, videos, and background images, and 3) detecting and scraping ad disclosure pages to determine whether an ad is targeted or not.

Detecting ads: To detect ads, we built on Zeng et al.’s [16] approach to use `EasyList`’s rules [109]. `EasyList` rules are commonly employed by popular adblockers to block or hide ads. For each detected ad element, the crawler recorded a set of attributes, including its position on the page, its dimensions, class, ID, and name, in addition to the complete HTML source and a screenshot. If the ad element contained any child elements, which was mostly true, the crawler recursively recorded their details, including all links, images, video, script, and iframe elements. Small elements

(< 30px in either dimension) and elements lacking any link, image, background image, or video were excluded.

The crawler not only took screenshots of each ad but also downloaded image and video elements that were descendants of the ad element. These media were utilized in the ML-based ad content analysis pipeline 4.6 alongside the ad screenshots. The crawler sent a single HTTP request during the page visit with the appropriate HTTP headers—such as the HTTP Referer [sic.] set to the current address-bar URL—when downloading these files. Finally, the crawler saved data-URL images found within the ads subtree. Bin Musa and Nithyanand [64] also utilized EasyList’s rules for ad identification in their study on tracker-advertiser relationships, but their implementation differs from ours. While they focus on detecting image-containing HTTP responses using the EasyList filter set, we query the DOM to detect ad elements, such as div elements, and their relevant descendant elements, such as images, iframes, links (a), and videos. Operating at the DOM level also allows us to detect and scrape ad disclosure pages to detect targeted ads. To verify how accurately our crawler detects ads, we performed a sanity check on a random sample of 105 ads (15 ads from each crawl). The crawler correctly detected ads in 85% of cases, misidentified non-ads in 7.5%, and captured blank or empty ads in 7.5%. Some ad screenshots also included multiple (2.8%) or only part (4.5%) of the ads. However, the overall accuracy and quality of our ads appear to be higher than prior work by Zeng et al. [55], which reported 34% unrendered (blank/unreadable) ads. We attribute this difference in data quality to two potential reasons. First, we use a more realistic crawler equipped with anti-bot measures; and second, unlike Zeng et al., we opted not to click the ads—which may trigger more stringent anti-bot, anti-fraud protections that prevent the delivery or rendering of the ads. Further, to evaluate whether our EasyList-based detector missed any ads, we manually reviewed 50 random pages where no ads were detected by the crawler. Our review did not reveal any false negatives, suggesting that our ad detection was robust. We also verified the accuracy of the ad images separately downloaded by the crawler, finding all of them to be present in the ads shown on the page.

Determining targeting criteria: To measure the prevalence of targeted advertisements at scale, we automated the process of scraping ad disclosure (e.g., “Why this ad”) pages. While the content of ad disclosure pages may vary by ad platform, they generally explain in broad terms why a specific advertisement was shown to a user. The reasons may include, for instance, *Google’s estimation of your interests* or *Websites you’ve visited*. The disclosure pages may also contain information about the website and the advertiser, and whether ad targeting is turned off for the website or a specific ad. Two example disclosure pages for a targeted and non-targeted ad are shown in Figure 3.

Ad disclosure pages are reachable by clicking the AdChoices icon  and the “Why this ad” button for Google ads [110] and other ad providers. Initially, we attempted to detect the ad disclosure links using fuzzy image matching based on the AdChoices icon. However, we

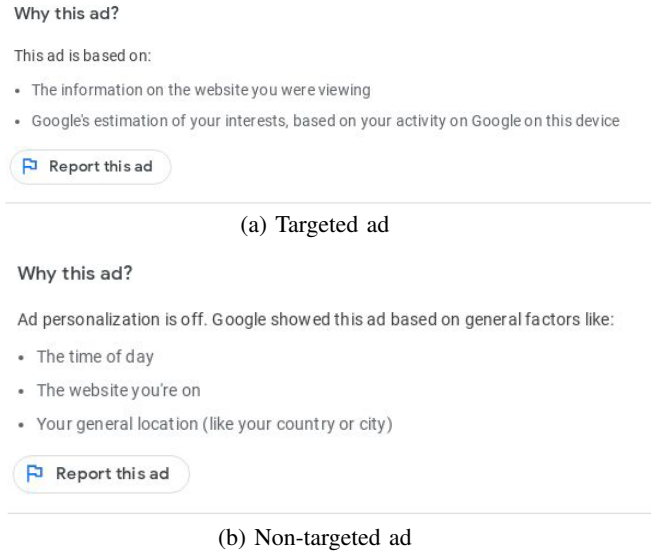


Figure 3: Google’s ad disclosure pages indicating whether an ad is targeted or not. The top figure belongs to a targeted ad (indicated by *Google’s estimation of your interests*), while the bottom one is for a non-targeted ad (indicated by *Ad personalization is turned off*)

found that the icon’s shape and visibility substantially vary across different ad vendors, and sometimes the icon can be hidden, making it unclickable. We chose to identify ad disclosure links through their URLs, focusing on a fixed set of providers we could detect reliably and deterministically. By analyzing ad disclosure pages in our pilot crawls, we compiled a list of hostnames (i.e., `adssettings.google.com`, `privacy.us.criteo.com` and `privacy.eu.criteo.com`) that appear in the ad disclosure links, and explain whether an ad is targeted or not. We limited our investigation to ad disclosure pages from these two providers because other providers we encountered in our pilot crawls did not offer useful information about the targeting criteria of the ads.

Once the crawler detects and clicks on the AdChoices link, the ad disclosure page opens in a new tab. We intercepted this new tab, stored its URL, screenshot and text contents (via `document.innerText`) for analysis. The scraped text contents are then used to detect whether ad targeting is enabled or not. Specifically, we searched in the ad disclosure texts for specific disclosure statements indicating whether and how an ad was targeted. The disclosure statements include, for instance, *Google’s estimation of your interests* (targeted), *Websites you’ve visited* (targeted) and *Ad personalization is turned off* (non-targeted). If one or more statements indicating targeted ads occur in an ad disclosure text, we label the ad as targeted. Otherwise, we label the ad as non-targeted. Note that we count behavioral or retargeted ads also in the targeted category. We compiled a list of 18 statements (Appendix A.2) incrementally, using over 40K ad disclosure texts extracted during the crawls. We ensured that

all ad disclosures contain at least one of these statements, to make sure our analysis is exhaustive.

Interacting with the page and ads: Upon loading a page, our crawler waited for 2.5 seconds and dismissed any popup dialogs using heuristics from prior work [29]. We dismissed these dialogs to prevent them from blocking our crawler’s interactions with the webpage. The crawler then waited for another second before scrolling through the page in 10 steps, taking strides of about 500–600 pixels each interlaced with a random delay of 500–600 milliseconds. Finally, after waiting for another second, it scrolled up to the beginning of the page using the same scrolling behavior. We engineered this up-and-down scrolling behavior to allow the webpage to load any ad slots that are lazily loaded as the user scrolls the page below the landing fold.

The crawler then identified all ads on the page. It set the border color of each ad to red to visually mark the ads for manual review. The crawler then took a screenshot of the entire page and then scraped each ad in a top-down fashion. To ensure that an advertisement is fully seen, it scrolled down to each ad before taking its screenshot. Finally, the crawler detected ad disclosure links and clicked each one individually to capture all ad disclosure texts and screenshots. We limited the number of scraped ads per page visit to ten, which limits over-representation by a few websites with many ads.

Analyzing advertisement content: We identified and measured four kinds of ads in our corpus: weight loss ads, mental health ads, dating services ads, and ads that contain clickbait racy content. While our dataset of ads can be used to perform fuller content analysis, we focused on these four categories since prior work [111], [112] and regulatory reports [113] have argued that these can be especially harmful to children. In fact, many ad networks’ moderation policies [114], [115] explicitly restrict these categories of ads from appearing on children’s websites. We note that the categories we focused on are not exhaustive, and our ads analysis is exploratory, serving as a preliminary investigation into this critical problem.

An overview of the ad content analysis pipeline is shown in Figure 4. To identify ads containing click-bait racy content, we employed Google Cloud Vision API’s SafeSearch Detection [116], which is a service that uses deep learning to analyze images and identify potentially unsafe content. It evaluates images against categories such as adult, violent, racy, and medical content and returns likelihood scores for each category, ranging from ‘VERY_UNLIKELY’ to ‘VERY_LIKELY.’ Upon manually evaluating the output generated by the algorithm, we focused on the ‘racy’ category with a likelihood of ‘VERY_LIKELY.’ We also tested Microsoft’s Adult Content Detection [117], part of Azure Cognitive Services, to identify racy images. However, due to more false positives compared to Google Cloud Vision API, we chose the latter for our study.

We used the Google Cloud Vision API to extract text from ad images following a similar approach to Bin Musa and Nithyanand [64]. The text in each image was extracted using the Optical Character Recognition (OCR) feature of

the API, specifically by employing the `fullTextAnnotation` attribute of the API response. This allowed us to extract text data at different levels, such as page, paragraph, and word. We opted to use the paragraph level since it gives the best separation in ads promoting multiple unrelated products. Despite their name, paragraphs returned by the API were relatively short and akin to sentences (21 characters, on average).

We then employed semantic similarity to identify the most similar ad texts (paragraphs) corresponding to a given search query, which in our case were “weight loss”, “mental health”, and “dating”. This approach is versatile and can be used to retrieve ads related to any arbitrary words or phrases. To compute the embeddings of the queries and ad paragraphs, we used the “paraphrase-multilingual-mpnet-base-v2” model, the distilled multilingual model we used to classify web pages (3.2). To find the most similar results, we calculated the cosine similarity between the embeddings of the search query and each ad paragraph and sorted them accordingly. Next, we manually reviewed the 100 most similar distinct paragraphs and their associated images, including ad screenshots or background ad images, to identify those that pertained to the three categories of interest. We also experimented with BERTopic [118] to create topic models and searched for clusters similar to our chosen categories. While this resulted in well-grouped texts, it required manual verification of numerous (several thousand) clusters. Sorting based on semantic similarity proved to be faster, more flexible, and easier to implement and evaluate, making it the preferred approach for manual reviewing.




Figure 4: Overview of the ad content analysis pipeline.

4.7. List of crawls

The main dataset used in our study consists of seven crawls (Table 2), all of which were run in April 2023 using cloud-based servers on Digital Ocean. The crawls include five desktop and two mobile crawls from five and two vantage points, respectively. We limited the mobile browser crawls to two vantage points because we do not focus on mobile-desktop comparison, which we leave to future work. The vantage points used for crawls consisted of Frankfurt, Amsterdam, London, San Francisco, and New York City (NYC), to capture ads and tracking from different jurisdictions. Crawls were run in parallel using separate servers

TABLE 2: Crawl statistics based on different vantage points. *: Avg./Sum of all sites visited across all crawls.

Form factor	Vantage point	Successfully loaded pages	Successful crawling rate
Desk.	NYC	10,310	95%
	SF	10,301	95%
	LON	10,270	95%
	FRA	10,221	95%
	AMS	10,014	93%
Mobile	NYC	10,168	94%
	FRA	10,283	96%
Avg./Sum		71,567	95%

with moderate resources (8 vCPU cores, 16GB RAM). Each crawl took between 13 and 32 hours to complete. A clean browser profile is used for each page visit to prevent ads targeted to our browsing history. During each crawl, we visited both landing and inner pages, following the process described in Section 4.3. When applicable, we accepted all personal data processing on consent (cookie) dialogs. The order of visited pages was randomized within each vantage point. Note that the San Francisco crawl used inner links extracted from the NYC crawl, while the London and Amsterdam crawl used inner links from the Frankfurt crawl. While this constraint did not appear to impact the success rate of visits across these vantage points, future research could explore identifying inner pages during the crawling process.

5. Measurement Results

Table 2 summarizes the overall statistics for measurement crawls. A total of 71,567 pages were loaded successfully across all crawls. The success rate of our crawler was over 93%, according to the successful visit criteria we developed and applied (Appendix A.3). For simplicity, certain comparative results presented below are based on desktop crawls from NYC and Frankfurt, representing one location each in the US and the EU.

5.1. Ad targeting and content analysis

Our crawler scraped 70,303 ads from 804 of the 2,004 distinct websites across seven crawls. An average of 36% of the pages contained one or more ads, and we detected targeted ads on 27% of the pages we crawled. The crawler scraped 10,839 and 9,447 ads on average in the crawls from the US and Europe, respectively.

5.1.1. Over 70% of ads with disclosures are targeted in nature. Our crawler captured a total of 40,281 ad disclosure pages, which we used to determine the advertiser’s identity and whether ad targeting is enabled or not. There are fewer disclosure pages than ads due to ads without disclosure links and failures in detecting or opening those links. In fact, we only consider ad disclosures from two ad providers: Google (97.8%) and Criteo (2.2%), since ad disclosure pages of other providers did not reveal the targeted status of the ad or




Figure 5: Tranco rank (x-axis) distribution of sites that use targeted vs. non-targeted ads. Popular websites (below) appear to be more prone to disabling ad targeting.

TABLE 3: Number of visits and scraped ads, along with percentages of ads/targeted ads per crawl. *: Percentage of targeted ads is only based on ads with disclosures. In the rightmost two columns, we include a site if we scraped at least one ad/targeted ad from one of its pages.

Form factor	Vantage point	# ads	% sites with ads	% sites with targeted ads	% targeted ads
Desk.	NYC	11,288	38%	30%	79%
	SF	10,950	38%	28%	67%
	LON	9,702	36%	27%	76%
	FRA	9,700	36%	26%	68%
	AMS	9,250	35%	26%	67%
Mobile	NYC	10,278	36%	29%	82%
	FRA	9,135	33%	26%	70%
Avg./Sum		70,303	36%	27%	73%

the advertisers’ identity. Limiting our analysis to 40,281 ads with disclosure pages, we found that targeting was enabled for 73% of the ads with disclosures. Comparing across different privacy jurisdictions, we find 68% of the ads on average were targeted in the EU crawls, compared to 76% in the UK and the US crawls. Comparing the crawls from the two US cities (SF & NYC), we find that 67% of the ads were targeted in the SF desktop crawl, compared to 79% and 82% in the NYC-based desktop and mobile crawls, respectively. Although these variations might be attributed to stricter privacy regulations like CCPA and GDPR, our available data and methods do not permit us to make this attribution. Comparing the Tranco ranks of the 689 websites that contain at least one targeted ad to 59 websites that only contain non-targeted ads, we find a tendency for popular websites to disable ad targeting (Figure 5). Sites with targeted ads had a median rank of $\sim 340K$, while those with only non-targeted ads had a median rank of $\sim 128K$. Note that we only include 40,281 ads for which we can determine the targeted status in this analysis.

5.1.2. Ads can be targeted from anywhere. The “About the advertiser” section in Google’s ad disclosures shows the name and location (country) of the advertisers. This information is only available in 70% of the ad disclosures in our dataset. Extracting these fields from the ad disclosure texts, we identified 1,685 distinct advertisers from 81 different countries. Advertisers with the most ads in our data are displayed in Table 4. We note that due to the transient, targeted and localized nature of ad campaigns, the list in

TABLE 4: Top ten advertisers by the number of ads across all crawls.

Advertiser	Location	# ads	% targeted	Type of ads
Vinden.nl B.V.	Netherlands	4,707	86%	Search results
EXPLORADS	Cyprus	3,265	73%	Search results
All Response	UK	2,453	68%	Search results
Glowworld LLC	USA	2,365	55%	Online learning
Amomama M.	Cyprus	921	72%	Workout muscle gain, weight loss
Media Quest	UAE	910	79%	Search results
Brain Metrics I.	Cyprus	814	50%	IQ tests
BetterMe	Cyprus	731	85%	Weight loss
Marketism	Israel	645	49%	Search results
Alibaba.com HK	Hong Kong	541	86%	Products sold on Alibaba.com

Table 4 may not represent the most common advertisers on child-directed websites in general. Further, in certain cases (e.g., Glowworld LLC and Marketism), an advertising or marketing agency is listed on the ad disclosure page instead of the company offering the advertised products or services.

The top ten advertisers are located in seven different countries and three continents. We observed that many of those advertisers are located far from our crawl vantage points, thus indicating that children visiting websites in our list can be targeted with ads from anywhere in the world. Reviewing a sample of 100 ads from each advertiser, we marked in the rightmost column their predominant ad theme. Five of the ten advertisers display ads for search results about various products—such as depression tests, belly fat reduction, senior meals, and electronic payments—on lesser-known search engines such as IngoSearch [119]. Ads from Betterme [120], a “behavioral healthcare app” with more than 100M installations, featured plans for weight loss, muscle gain, and intermittent fasting (e.g., Figure 1 (h)).² Brain Metrics Initiative displays ads for IQ tests, an example for which is given in Figure 1 (c). Alibaba Hong Kong, on the other hand, displays ads featuring racy and disturbing images of products sold on alibaba.com. For instance, the ad on the top left (a) in Figure 1 features recurring images in Alibaba ads: a naked baby model (leftmost), rabbit meat (rightmost), and a semi-transparent underwear ad in the middle. We investigate similar racy clickbait ads and other improper ads in the following subsection.

5.1.3. Improper ads on child-directed sites. In total, our crawler collected 199,935 screenshots and images from the 70,303 scraped ads. After deduplicating the images, we queried the Cloud Vision API to obtain the category and OCR texts of the resulting 98,264 distinct images. We manually reviewed 741 images classified as ‘VERY_LIKELY’ racy by the API. Separately, we reviewed 1,136 ad images with OCR text semantically most similar to our search terms (mental health, dating, and weight loss). Due to study

2. We note that Better.me’s data sharing practices with third parties were investigated by Privacy International, but the company reportedly took corrective action[121].

TABLE 5: Number of improper ads identified for each crawl.

Form factor	Vantage point	Dating	Mental health	Weight loss	Racy	Some-what racy	Total
Desk.	NYC	4	21	16	21	26	88
	SF	7	9	15	6	25	62
	LON	10	17	48	12	31	118
	FRA	1	0	48	19	25	93
	AMS	8	4	82	10	33	137
Mobile	NYC	22	25	113	98	17	275
	FRA	18	5	190	11	6	230
Total		70	81	512	177	163	1003

limitations, we only examined the ads related to the top 100 distinct texts for each term. Since each distinct text may appear in multiple ads differently, we labeled the images separately and used videos captured by the crawler when the ad was animated or the ad screenshot was obscured. Table 5 shows the number of improper ads identified in each crawl, amounting to 1,003 across 311 distinct websites. A notable finding is the higher prevalence of such ads on mobile devices compared to desktops in general.

Racy images. We found 177 racy ads and 163 somewhat racy ads, considered edge cases due to their potential inappropriateness for child-directed websites. These ads were identified across 80 distinct websites mostly ranked within the top one million according to the Tranco list, with a median rank of 426K. Figure 1 (a), (g), (k) are examples of some of these ads. Notably, the majority of the racy ads were encountered in mobile crawls; especially within the NYC crawl (98/177). Out of 177 racy ads, only 38 had disclosure pages. Among these, targeting was enabled for 35 ads.

Mental health. By manually labeling 236 ad images, we identified 81 ads offering mental health services on 48 distinct websites. Examples of ads in this category contained “take a depression test” (Figure 1 (f)), “online psychiatrists,” “how to get over depression,” and a “mental health chatbot which helps people with depression.” We excluded false positives that were not mental health service offerings, such as ads for “mental health counselor salaries,” “online psychology courses,” and “psychology books.”

Dating. Manually labeling 231 ad images, we identified 70 dating platform ads on 48 distinct websites, most of which targeted mobile users. The ads promoted dating platforms such as “dating.com,” and “Live Me,” a live streaming app with ads featuring suggestive imagery (Figure 1, (j), (l)). Another ad for DateMyAge.com featured a call to “[m]eet your mature singles” (e). False positives removed during the manual labeling of this category included ads for customer relationship tools, romantic holiday tours, and online appointment services.

Weight loss. We identified 512 weight loss-related ads (plans, apps, products) on 170 distinct websites by labeling 669 ad images. Notably, there was a higher number of weight loss ads on mobile devices, indicating campaigns targeting mobile users. Examples of text featured in these ads

TABLE 6: Average number of third-party and tracker domains, and the prevalence of tracking and fingerprinting on child-directed websites from five vantage points.

Form factor	Vantage point	3rd-Party domains	Tracker domains	Tracker entities	% sites with 3rd Parties	% sites with trackers	% site with FP
Desk.	NYC	31.6	23.4	20.0	95%	90%	9%
	SF	29.3	21.3	17.8	95%	91%	9%
	LON	21.3	14.3	10.6	96%	91%	7%
	FRA	23.2	15.6	11.7	95%	90%	10%
	AMS	21.4	14.3	10.6	93%	89%	7%
Mobile	NYC	29.8	21.8	18.4	95%	91%	9%
	FRA	22.6	15.2	11.5	95%	90%	11%

included “intermittent fasting for weight loss,” “keto weight loss plan,” and “eating plan to lose weight” (Figure 1 (h)).

In Figure 1, we provide additional examples of advertisements that are likely not suitable for children. Examples of these included an ad for a test called “Am I Gay Test” (d), for a sex toy (e) and a sex toy shop (f)³ featuring an image of ice cream that could be appealing to children, and ads featuring clickbait and sexually suggestive images. The ads were found on websites related to K-12 e-learning, kids games, coloring books and worksheets, among others. The ads in Figure 1 do not necessarily fit our four investigated categories but showcase the diversity of improper ads.

Malicious ad links. Finally, we present an exploratory analysis of whether ads on child-directed websites link to malicious pages. We submitted a sample of links extracted from the ad elements to the VirusTotal API in August 2023. Specifically, we removed links with duplicate hostnames, and for Google ads, we extracted a direct link to the ad landing page using the ‘adurl’ parameter [123]. While the overwhelming majority of the links were classified as benign, 149 of the nearly 3,940 scanned links were flagged as malicious or phishing by at least one scan engine. Notably, the word “taboola” was mentioned in 78 of the 149 detected links as a URL parameter that seems to indicate the ad network (`network=taboola`).

5.2. Tracking and fingerprinting analysis

Table 6 shows the prevalence of third-party trackers detected across different crawls. We find that around 90% of the websites have at least one tracker domain, and over 93% embed at least one third-party domain.

Third-party trackers. The average number of tracker domains per site differs significantly, e.g., 15.6 and 23.4 in Frankfurt and NYC crawls, respectively, while the median is 15 and 16, respectively. The difference in averages can be attributed to websites with the high number of trackers in the NYC crawl. This explanation is in line with the results displayed in Table 8, which shows the top five websites with the most trackers in Frankfurt and NYC crawls. Most of these websites are among the top one million, which means they likely receive substantial traffic. Notably, all of these sites displayed ads that were targeted. The numbers shown in

3. Reportedly Germany’s largest online adult retailer [122]

TABLE 7: Prevalence of tracker entities in terms of number of distinct websites in Frankfurt and NYC desktop crawls.

FRA		NYC	
Entity	# Sites	Entity	# Sites
Google	1,702	Google	1,718
Facebook	458	Microsoft	549
Index Exchange	424	Adobe	543
Xandr	416	Xandr	516
Adform	412	The Trade Desk	501
The Trade Desk	390	Index Exchange	495
OpenX	378	IPONWEB	467
Adobe	366	Facebook	456
Quantcast	361	Magnite	446
PubMatic	359	OpenX	426

the table - number of trackers, requests, and cookies - reflect averages across the web pages. In the NYC crawl, visiting `mathfunworksheets.com` triggered a total of 1,547 requests involving 161 unique third-party tracker entities (i.e., organizations/companies). Another website, `woojr.com` found to contain 148 distinct third-party tracker entities when visited from NYC. This website includes resources for children’s activities and educational materials, including printable worksheets and fun activity pages. When visited from Frankfurt, `www.wowescape.com`, a website offering various games for children and teenagers, triggered requests to 95 distinct third-party tracker entities.

Most prevalent trackers. Table 7 shows the tracker entities with the most prevalence in Frankfurt and NYC desktop crawls. We found a tracking-related request to Google domains, including its analytics, advertising and tag management scripts on ~84% of the 2,004 child-directed websites in both crawls. Facebook is the second most prevalent entity in the Frankfurt crawl mostly due to Facebook Pixel (on 427 websites), which facilitates ad retargeting and conversion measurement, among others [124]. Largely thanks to Linked Insight Tag (`px.ads.linkedin.com`, 466 websites), Microsoft is the second most prevalent entity in the NYC crawl. Linked Insight Tag serves multiple purposes, including retargeting, conversion measurement, and providing demographic insights about website visitors [125].

Regional differences. To explore the differences in tracker entities across vantage points, we compared the tracker entities from Frankfurt and NYC desktop crawls. Despite a considerable overlap among the detected tracker entities (Jaccard index=0.85), we also identified variations. Specifically, our investigation unveiled 47 tracker entities exclusive to the Frankfurt crawl and 118 tracker entities that were only found in the NYC crawl. For instance, tracking related requests to *advanced STORE* [126] (`ad4m.at` & `ad4mat.net`, 236 websites) exclusively appear in the crawl from Frankfurt, whereas *Throtle*, a company that provides an identity graph to marketers and advertisers, only appears on 171 websites in the NYC crawl [127]. Furthermore, we find that the majority of the websites in both Frankfurt and NYC crawls (70% and 72%, respectively) contain third-party trackers that set at least one cookie with the `SameSite=None` attribute and a lifespan of over three months. Primarily through `doubleclick.net` domain,




Figure 6: Comparison of the average number of third-party and tracker domains/entities on websites containing ads vs. not containing ads.

TABLE 8: websites with the most distinct tracker entities. The table shows the websites’ distinct third-party tracker entities, the number of requests, cookies, and Tranco rank.

Loc.	Website	# Trackers	# Requests	# Cookies	Rank
NYC	mathfunworksheets.com	161	1,547	395	669K
	woojr.com	148	2,181	391	83K
	innerchildfun.com	139	1,235	336	308K
	kidzfeed.com	138	1,050	272	797K
	thecolor.com	138	1,068	260	192K
FRA	wowescape.com	95	392	55	258K
	webgames.io	94	564	92	155K
	coloriages-pour-enfants.net	90	401	66	919K
	theschoolrun.com	87	417	91	760K
	testsworld.net	86	478	138	-

Google set these cookies on over 51% of the websites. While identifying the individual purposes of these cookies is out of scope, this combination of cookie attributes (esp. setting SameSite=None) makes it possible to track users across websites.

Sites with and without ads. As part of our investigation, we conducted an additional analysis to compare how the number of third parties and trackers change between websites with and without ads. Figure 6 shows that websites with ads tend to have substantially more third-party and tracker domains. More specifically, the figure shows websites with ads tend to contain two to four times more third-party and tracker domains.

Browser Fingerprinting. We now discuss our findings on fingerprinting scripts on child-directed websites. Table 6 shows that we detect fingerprinting scripts on 176 (9%) and 218 (10%) websites in Frankfurt and NYC crawls, respectively. The overall prevalence of fingerprinting aligns with the recent research by Iqbal et al., which finds fingerprinting on 10.18% of the top-100K websites [38]. One of the most prevalent fingerprinters in both crawls is an online payment company (Stripe; 66, 67 sites on Frankfurt and NYC crawls, respectively). According to their help pages [128] Stripe primarily employs fingerprinting for fraud prevention purposes. Webgains (82 sites in the Frankfurt crawl), an affiliate marketing company, also mentions fingerprinting in their Data Processing Agreement with Merchants [129], but without specifying its purpose. The most commonly used fingerprinting method is Canvas fingerprinting, present on about 208 sites in the Frankfurt crawl and about 172 sites

in the NYC crawl.

We found one or more trackers to be present on more than 90% of mobile websites (Table 6), which is similar to our finding for the desktop websites. NYC and Frankfurt crawls differ slightly in the number of ads: we scraped 10,278 ads in the NYC crawl and only 9,135 in the Frankfurt crawl—the latter is the crawl with the least amount of ads. Slightly more (29 vs 26%) websites in the NYC mobile crawl have targeted ads; and NYC mobile crawl has the highest proportion of targeted ads (82%) across all crawls. We also discovered that improper ads, particularly racy and weight loss ads, were more prevalent on mobile devices compared to desktops.

6. Discussion

Our research paints a troubling picture of tracking and inappropriate advertising practices on child-directed websites. Advertisements featuring sexually suggestive imagery and ads about weight loss, dating, and mental health may pose potential risks to children’s emotional and psychological welfare. We discuss the legal implications, ethical considerations and limitations of our study below.

6.1. Legal implications

In this section, we discuss what the law says about tracking and advertising practices uncovered in our research. We focus on the EU General Data Protection Regulation (GDPR) and the US Children’s Online Privacy Protection Act (COPPA)⁴.

The GDPR and the ePrivacy Directive. Under the GDPR, companies are only allowed to process personal data if they have a ‘legal basis’ for such processing. The GDPR provides six possible legal bases (article 6 GDPR). However, generally, the data subject’s consent is the only possible legal basis for online tracking and behavioral (targeted) advertising [130]. Moreover, the ePrivacy Directive [131] requires, in short, companies to ask the internet user for consent before they use tracking cookies or similar tracking technologies (Article 5(3)).

The GDPR’s requirements for valid consent are strict. Consent is only valid if it is really voluntary (‘freely given’), and ‘specific’ and ‘informed’. The data controllers (the website owner and the companies involved in tracking and targeted advertising) must ‘be able to demonstrate that the data subject has consented to process of his or her personal data’ (Article 7(1) GDPR). The GDPR’s requirements for valid consent also apply to consent (for cookies etc.) as prescribed by the ePrivacy Directive. The GDPR has specific rules for consent by children. Roughly summarized, children cannot give valid consent; the parent should give consent instead (Article 8 GDPR). EU member states have set different minimum consent ages, ranging from 13 to 16 years [132].

4. We do not analyze whether specific companies breach the law. For such an analysis, each case would have to be examined separately, considering all the circumstances of that specific case. Rather, we discuss legal requirements in general terms.

Hence, only parental consent can legitimize tracking on a children’s website. Observe that a parent clicking a consent dialog (as done by our crawler) does not constitute parental consent under GDPR. Even in low-risk cases, verification of parental responsibility via email may be necessary [133].

The EU Digital Services Act. The rules for tracking and targeting children will become stricter in the EU. From 17 February 2024 on, the EU Digital Services Act [134] applies. Article 28 says, roughly summarized, that online platforms must not use behavioral advertising ‘when they are aware with reasonable certainty that the recipient of the service is a minor’ [134]. This prohibition cannot be overridden with the consent of the child or the parent. The DSA also requires “very large online platforms” [135] (with more than 45 million users in the EU) to publish the advertisements that it presented to users in an online repository, together with information about, for instance, the targeting criteria (Article 33, 39 DSA). The methods that we used in this paper could be used to check the completeness and accuracy of data published in those repositories.

COPPA. COPPA regulates companies offering a website or online service directed to children under the age of 13. Specifically, COPPA applies to companies using children’s ‘personal information,’ which includes ‘persistent identifiers such as cookies and device fingerprints’ (COPPA §312.2) [136]. The website owner is responsible for data collection by third parties through its site. Such third parties must also comply with COPPA. Companies based outside the US must also comply with COPPA if their services are directed to children in the US [136].

Our results showed that 27% of the child-directed websites use targeted advertising. Under COPPA, data collection for targeted advertising on these websites is only allowed after getting parents Verifiable Parental Consent (VPC). VPC entails utilizing stringent verification methods, such as credit card verification, face recognition, or government ID checks [137]. This makes VPC much more complex than simply clicking an accept button on a dialog. We note that our crawler cannot simply give VPC.

6.2. Research Ethics

Our crawler visited over 166K pages and it triggered many ad impressions that could be viewed by a real visitor (likely a child). Given the huge scale of the digital ad market (projected to reach US\$700bn in 2023 [138]) we believe these ad impressions are a negligible cost for raising the transparency around tracking and ads targeted to children. Furthermore, we took several measures to limit our footprint on the crawled websites. For instance, we only crawled five inner pages from each site in a crawl, and we randomly shuffled the target URLs to avoid concurrently visiting the inner pages of a website. We also took appropriate measures to ensure that no harm was done to collaborators involved in the project, especially when dealing with explicitly graphic images.

Disclosures and outreach In July 2023, we reached out to five companies that we found to serve racy ads.

One company invited us to a call and explained the likely reason for showing inappropriate ads (websites failing to label themselves as child-directed and a false negative in the ad company’s automated child-directed site detection tool). They added the sites in question to the child-directed list and pledged further investigation. Another thanked us and began an internal review. The third redirected our query to the relevant department. Moreover, we disclosed 34 racy ads to Google by manually visiting the ad disclosure URLs of each racy (Google) ad; and using the *Report this ad* button. Note that, to identify the ad vendors involved in serving the ad, we used a combination of ad images, and src/href attributes of the ads’ descendant iframe, image and link elements (§4.6).

In addition, we shared our preliminary results with a European data protection agency (DPA), and a consumer protection agency. Both showed interest; the DPA asked if there were any websites from their country containing improper ads. The consumer protection agency invited us to present our work. We also shared our findings with civil society and industry organizations including the 5Rights Foundation [139]. We plan to further share our study with regulators and other relevant stakeholders.

While using the VirusTotal API, we found and reported three porn websites miscategorized as kids-related to the respective third-party categorization service. Although no response was received, the categories were later rectified.

6.3. Limitations

Our study predominantly covers websites targeting younger children, as we define ‘children’ as under 13, aligning with both US and EU regulations. While our classifier detected child-directed sites in 48 languages, it may have a potential bias towards English content, which is over-represented in the training data. The classifier may prefer sites with descriptive titles and content or may carry biases related to website age, design, or accessibility. Moreover, the classifier favors precision over recall to reduce the manual labeling workload.

Our research is the first attempt to build a large list of child-directed websites. Classifying websites as child-directed or not is challenging, primarily due to the existence of gray areas that complicate the labeling process. For instance, many websites offer content that appeals to both children and adults. While we tried to exclude websites that can only be used by teachers or parents, certain websites included pages such as online exercises, videos or games that can be consumed by children. To validate our list, we performed further analysis on two subsets of websites: two senior researchers relabeled a random 100-website sample; one of the senior researchers relabeled a random subset of 100 websites with targeted (50) and inappropriate ads (50). In both cases, we identified less than 9% of websites that are related to children, but mainly catered to teachers or parents (8.6%, 8.3%, respectively). While this impurity could be avoided with more conservative labeling, our analysis of 200 websites strongly suggests these cases do not skew our primary findings.

While we found fewer targeted ads in the EU than in the US, we cannot directly attribute this to differences in privacy regulation or another specific factor. Failure to detect and interact with consent dialogs may be a confounding factor, among others.

When detecting targeted ads, we only used ad disclosure pages from two providers (Google and Criteo) due to the unavailability of useful ad disclosures from other vendors. Thus, our targeted ad detection method depends on the accuracy, completeness and precision of Google and Criteo's ad disclosures.

Websites may treat cloud-based IP addresses or automated browsers differently [140], [141], [39]. To curb such effects, we used the anti-bot detection features of TRC [93]. Reviewing the screenshots captured during the visits, we observed very few blocked visits.

We conducted four sets of inner link collection crawls: two from NYC and two from Frankfurt, encompassing both desktop and mobile crawls. This constraint does not appear to impact the success rate of visits across these vantage points; nonetheless, future research could explore the possibility of identifying inner pages during the crawling process.

Since we use a fresh profile for each page visit, we may not capture re-targeted or other personalized ads that are only shown to users with a behavioral profile. Future work could extend our method to incorporate personas and warm-up crawls to study such ads. Overall we do not claim that our findings are representative of tracking and advertising practices on child-directed websites. Our focus in this study is not on how ads are targeted, but simply on *whether the targeting is enabled or not*.

7. Conclusion

We presented an empirical study of online tracking and advertisements on over 2,000 child-directed websites. Building a lightweight and versatile ML pipeline to analyze ad content, we identify hundreds of cases of improper ads, including weight loss and mental health ads, and ads featuring dating services, racy and sexually suggestive imagery. Our study reveals several notable trends: websites featuring advertisements tend to contain two to four times more trackers, mobile websites exhibit a greater prevalence of inappropriate ads, and popular websites are less likely to deploy targeted advertisements. Our findings provide concrete evidence of troublesome practices that are likely illegal, unethical, or simply careless. We call for more research, regulation and enforcement to limit the ongoing violation of children's privacy and well-being.

8. Acknowledgments

Asuman Senol was funded by the Cyber-Defence (CYD) Campus of armasuisse Science and Technology. Veelasha Moonsamy was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germanys Excellence Strategy - EXC 2092 CASA - 390781972 and the BMBF project DigiFit - grant number 16KIS1646K.

We thank our shepherd and the reviewers for their constructive feedback.

References

- [1] S. Englehardt and A. Narayanan, "Online tracking: A 1-million-site measurement and analysis," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 2016, pp. 1388–1401.
- [2] M. D. Ayenson, D. J. Wambach, A. Soltani, N. Good, and C. J. Hoofnagle, "Flash Cookies and Privacy II: Now with HTML5 and ETag Respawning," 2011.
- [3] I. Fouad, N. Bielova, A. Legout, and N. Sarafijanovic-Djukic, "Missed by filter lists: Detecting unknown third-party trackers with invisible pixels," *Proceedings on Privacy Enhancing Technologies*, vol. 2020, no. 2, pp. 499–518, 2020.
- [4] G. Acar, C. Eubank, S. Englehardt, M. Juarez, A. Narayanan, and C. Diaz, "The web never forgets: Persistent tracking mechanisms in the wild," in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, 2014, pp. 674–689.
- [5] A. Zarras, A. Kapravelos, G. Stringhini, T. Holz, C. Kruegel, and G. Vigna, "The dark alleys of madison avenue: Understanding malicious advertisements," in *Proceedings of the Conference on Internet Measurement Conference*, 2014, pp. 373–380.
- [6] Z. Li, K. Zhang, Y. Xie, F. Yu, and X. Wang, "Knowing your enemy: understanding and detecting malicious web advertising," in *Proceedings of the ACM conference on Computer and communications security*, 2012, pp. 674–686.
- [7] K. Subramani, X. Yuan, O. Setayeshfar, P. Vadrevu, K. H. Lee, and R. Perdisci, "When push comes to ads: Measuring the rise of (malicious) push advertising," in *Proceedings of the ACM Internet Measurement Conference*, 2020, pp. 724–737.
- [8] I. Ilascu, "Hackers push malware via Google search ads for VLC, 7-Zip, CCleaner," *BleepingComputer*, Jan. 2023, <https://www.bleepingcomputer.com/news/security/hackers-push-malware-via-google-search-ads-for-vlc-7-zip-ccleaner>.
- [9] L. Abrams, "Ransomware access brokers use Google ads to breach your network," *BleepingComputer*, Jan. 2023, <https://www.bleepingcomputer.com/news/security/ransomware-access-brokers-use-google-ads-to-breach-your-network>.
- [10] "MalVirt | .NET Virtualization Thrives in Malvertising Attacks," <https://www.sentinelone.com/labs/malvirt-net-virtualization-thrives-in-malvertising-attacks>, 2023, [Accessed 28 Feb. 2023].
- [11] "Now even YouTube serves ads with CPU-draining cryptocurrency miners," <https://arstechnica.com/information-technology/2018/01/now-even-youtube-serves-ads-with-cpu-draining-cryptocurrency-miners/>, 2018, [Accessed 28 Feb. 2023].
- [12] "Rogue GIMP Google Ad Pushed Info-Stealer Malware Through Website Replica," <https://www.bitdefender.com/blog/hotforsecurity/rogue-gimp-google-ad-pushed-info-stealer-malware-through-website-replica/>, 2022, [Accessed 28 Feb. 2023].
- [13] E. Tekiner, A. Acar, A. S. Uluagac, E. Kirda, and A. A. Selcuk, "SoK: cryptojacking malware," in *2021 IEEE European Symposium on Security and Privacy (EuroS&P)*, pp. 120–139.
- [14] K. C. Wilbur and Y. Zhu, "Click fraud," *Marketing Science*, vol. 28, no. 2, pp. 293–308, 2009.
- [15] "Estimated cost of digital ad fraud worldwide from 2018 to 2023," <https://www.statista.com/statistics/677466/digital-ad-fraud-cost/>, 2023, [Accessed 28 Feb. 2023].
- [16] E. Zeng, T. Kohno, and F. Roesner, "What makes a "bad" ad? user perceptions of problematic online advertising," in *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 2021, pp. 1–24.

- [17] J. Zhao, G. Wang, C. Dally, P. Slovak, J. Edbrooke-Childs, M. Van Kleek, and N. Shadbolt, "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*, New York, NY, USA, 2019, p. 113.
- [18] P. Kumar, S. M. Naik, U. R. Devkar, M. Chetty, T. L. Clegg, and J. Vitak, "'No Telling Passcodes Out Because They're Private': Understanding Children's Mental Models of Privacy and Security Online," *Proc. ACM Hum.-Comput. Interact.*, vol. 1, no. CSCW, dec 2017.
- [19] M.-A. Lawlor and A. Prothero, "Pester power—a battle of wills between children and their parents," *Journal of Marketing Management*, vol. 27, no. 5-6, pp. 561–581, 2011.
- [20] D. Kunkel, B. L. Wilcox, J. Cantor, E. Palmer, S. Linn, and P. Dowrick, "Report of the APA task force on advertising and children," *Washington, DC: American Psychological Association*, vol. 30, p. 60, 2004.
- [21] D. R. John, "Consumer socialization of children: A retrospective look at twenty-five years of research," *Journal of consumer research*, vol. 26, no. 3, pp. 183–213, 1999.
- [22] M. Buijzen, "Reducing children's susceptibility to commercials: Mechanisms of factual and evaluative advertising interventions," *Media Psychology*, vol. 9, no. 2, pp. 411–430, 2007.
- [23] X. Cai and X. Zhao, "Online advertising on popular childrens websites: Structural features and privacy issues," *Computers in Human Behavior*, vol. 29, no. 4, pp. 1510–1518, 2013.
- [24] E. Rozendaal, M. Buijzen, and P. Valkenburg, "Comparing children's and adults' cognitive advertising competences in the Netherlands," *Journal of Children and Media*, vol. 4, no. 1, pp. 77–89, 2010.
- [25] M. Ali, M. Blades, C. Oates, and F. Blumberg, "Young children's ability to recognize advertisements in web page designs," *British Journal of Developmental Psychology*, vol. 27, no. 1, pp. 71–83, 2009.
- [26] V. Le Pochat, T. Van Goethem, S. Tajalizadehkhoob, M. Korczyński, and W. Joosen, "Tranco: A Research-Oriented Top Sites Ranking Hardened Against Manipulation," in *Proceedings of the 26th Annual Network and Distributed System Security Symposium*, 2019.
- [27] "Alexa.com (from the Wayback Machine Archive)," <https://web.archive.org/web/20210327233130/https://www.alexa.com>, 2023, [Accessed 28 Feb. 2023].
- [28] "We retired Alexa.com on May 1, 2022," <https://web.archive.org/web/20221020130106/https://support.alexa.com/hc/en-us/articles/4410503838999>, 2023, [Accessed 28 Feb. 2023].
- [29] A. Mathur, G. Acar, M. J. Friedman, E. Lucherini, J. Mayer, M. Chetty, and A. Narayanan, "Dark patterns at scale: Findings from a crawl of 11k shopping websites," *Proceedings of the ACM on Human-Computer Interaction*, vol. 3, no. CSCW, pp. 1–32, 2019.
- [30] P. Vallina, V. Le Pochat, . Feal, M. Paraschiv, J. Gamba, T. Burke, O. Hohlfeld, J. Tapiador, and N. Vallina-Rodriguez, "Mis-shapes, Mistakes, Misfits: An Analysis of Domain Classification Services," in *Internet Measurement Conference*, ser. IMC 2020, pp. 598–618.
- [31] N. Vlajic, M. El Masri, G. M. Riva, M. Barry, and D. Doran, "Online tracking of kids and teens by means of invisible images: COPPA vs. GDPR," in *Proceedings of the 2nd International Workshop on Multimedia Privacy and Security*, 2018, pp. 96–103.
- [32] "June/July 2022 crawl archive now available – Common Crawl," <https://commoncrawl.org/2022/07/june-july-2022-crawl-archive-now-available>, 2023, [Accessed 28 Feb. 2023].
- [33] "Digital Services Package: Commission welcomes the adoption by the European Parliament of the EU's new rulebook for digital services," https://ec.europa.eu/commission/presscorner/detail/en/IP_22_4313, 2022, [Accessed 28 Feb. 2023].
- [34] "State of the Union Address," <https://www.whitehouse.gov/state-of-the-union-2023/>, 2023, [Accessed 28 Feb. 2023].
- [35] "Senators Markey and Cassidy Reintroduce COPPA 2.0, Bipartisan Legislation to Protect Online Privacy of Children and Teens," <https://www.markey.senate.gov/news/press-releases/senators-markey-and-cassidy-reintroduce-coppa-20-bipartisan-legislation-to-protect-online-privacy-of-children-and-teens>, 2023, [Accessed 28 Feb. 2023].
- [36] FTC, "Children's Online Privacy Protection Rule ("COPPA")," <https://www.ftc.gov/node/60175>, 2023, [Accessed 28 Feb. 2023].
- [37] F. Roesner, T. Kohno, and D. Wetherall, "Detecting and defending against third-party tracking on the web," in *Presented as part of the 9th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 12)*, 2012, pp. 155–168.
- [38] U. Iqbal, S. Englehardt, and Z. Shafiq, "Fingerprinting the fingerprinters: Learning to detect browser fingerprinting behaviors," in *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2021, pp. 1143–1161.
- [39] D. Cassel, S.-C. Lin, A. Buraggina, W. Wang, A. Zhang, L. Bauer, H.-C. Hsiao, L. Jia, and T. Libert, "OmniCrawl: Comprehensive Measurement of Web Tracking With Real Desktop and Mobile Browsers," *Proc. Priv. Enhancing Technol.*, vol. 2022, no. 1, pp. 227–252, 2022.
- [40] S. Dambra, I. Sanchez-Rola, L. Bilge, and D. Balzarotti, "When Sally Met Trackers: Web Tracking From the Users' Perspective," in *31st USENIX Security Symposium (USENIX Security 22)*, 2022, pp. 2189–2206.
- [41] I. Sanchez-Rola, X. Ugarte-Pedrero, I. Santos, and P. G. Bringas, "The web is watching you: A comprehensive review of web-tracking techniques and countermeasures," *Logic Journal of the IGPL*, vol. 25, no. 1, pp. 18–29, 2017.
- [42] D. Kristol and L. Montulli, "RFC2965 - HTTP state management mechanism," Tech. Rep., 2000.
- [43] Q. Chen, P. Ilia, M. Polychronakis, and A. Kapravelos, "Cookie swap party: Abusing first-party cookies for web tracking," in *Proceedings of the Web Conference*, 2021, pp. 2117–2129.
- [44] J. R. Mayer and J. C. Mitchell, "Third-party web tracking: Policy and technology," in *2012 IEEE symposium on security and privacy*. IEEE, 2012, pp. 413–427.
- [45] P. Eckersley, "How unique is your web browser?" in *Privacy Enhancing Technologies: 10th International Symposium, Berlin, Germany, July 21-23, 2010. Proceedings 10*. Springer, 2010, pp. 1–18.
- [46] T. Kohno, A. Broido, and K. C. Claffy, "Remote physical device fingerprinting," *IEEE Transactions on Dependable and Secure Computing*, vol. 2, no. 2, pp. 93–108, 2005.
- [47] K. Mowery and H. Shacham, "Pixel perfect: Fingerprinting canvas in HTML5," *Proceedings of W2SP*, 2012.
- [48] Y. Cao, S. Li, and E. Wijmans, "(Cross-) browser fingerprinting via os and hardware level features," in *Proceedings Network and Distributed System Security Symposium*. Internet Society, 2017.
- [49] T. Laor, N. Mehanna, A. Durey, V. Dyadyuk, P. Laperdrix, C. Maurice, Y. Oren, R. Rouvoy, W. Rudametkin, and Y. Yarom, "DRAW-NAPART: A Device Identification Technique based on Remote GPU Fingerprinting," in *Network and Distributed System Security Symposium (NDSS)*, 2022.
- [50] P. Laperdrix, O. Starov, Q. Chen, A. Kapravelos, and N. Nikiforakis, "Fingerprinting in Style: Detecting Browser Extensions via Injected Style Sheets," in *USENIX Security Symposium*, 2021, pp. 2507–2524.
- [51] A. Sjösten, D. Hedin, and A. Sabelfeld, "Essentialfp: Exposing the essence of browser fingerprinting," in *IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, 2021, pp. 32–48.
- [52] U. Iqbal, P. Snyder, S. Zhu, B. Livshits, Z. Qian, and Z. Shafiq, "Adgraph: A graph-based approach to ad and tracker blocking," in *IEEE Symposium on Security and Privacy (SP)*, 2020, pp. 763–776.

- [53] S. Siby, U. Iqbal, S. Englehardt, Z. Shafiq, and C. Troncoso, “WebGraph: Capturing advertising and tracking information flows for robust blocking,” in *31st USENIX Security Symposium (USENIX Security 22)*, 2022, pp. 2875–2892.
- [54] I. Reyes, P. Wijesekera, J. Reardon, A. Elazari Bar On, A. Razaghpahanah, N. Vallina-Rodriguez, S. Egelman *et al.*, ““Won’t somebody think of the children?” examining COPPA compliance at scale,” in *The 18th Privacy Enhancing Technologies Symposium*, 2018.
- [55] E. Zeng, T. Kohno, and F. Roesner, “Bad news: Clickbait and deceptive ads on news and misinformation websites,” in *Workshop on Technology and Consumer Protection*, 2020, pp. 1–11.
- [56] E. Zeng, M. Wei, T. Gregersen, T. Kohno, and F. Roesner, “Polls, clickbait, and commemorative \$2 bills: problematic political advertising on news and media websites around the 2020 US elections,” in *Proceedings of the 21st ACM Internet Measurement Conference*, 2021, pp. 507–525.
- [57] O. Akgul, R. Roberts, M. Namara, D. Levin, and M. L. Mazurek, “Investigating Influencer VPN Ads on YouTube,” in *IEEE Symposium on Security and Privacy (SP)*, 2022, pp. 876–892.
- [58] M. Ali, A. Goetzen, A. Mislove, E. Redmiles, and P. Sapiezynski, “All things unequal: Measuring disparity of potentially harmful ads on Facebook,” in *6th Workshop on Technology and Consumer Protection*, 2022.
- [59] G. Venkatadri, E. Lucherini, P. Sapiezynski, and A. Mislove, “Investigating sources of PII used in Facebook’s targeted advertising,” *Proc. Priv. Enhancing Technol.*, vol. 2019, no. 1, pp. 227–244, 2019.
- [60] Y. Zhao, T. Liu, H. Wang, Y. Liu, J. Grundy, and L. Li, “Are mobile advertisements in compliance with apps age group?” in *Proceedings of the ACM Web Conference 2023*, 2023, pp. 3132–3141.
- [61] T. Medjkoune, O. Goga, and J. Senechal, “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*, ser. CCS ’23. New York, NY, USA: Association for Computing Machinery, 2023, p. 180194. [Online]. Available: <https://doi.org/10.1145/3576915.3623172>
- [62] A. Andreou, G. Venkatadri, O. Goga, K. P. Gummedi, P. Loiseau, and A. Mislove, “Investigating ad transparency mechanisms in social media: A case study of Facebook’s explanations,” in *NDSS -Network and Distributed System Security Symposium*, 2018, pp. 1–15.
- [63] M. Eslami, S. R. Krishna Kumaran, C. Sandvig, and K. Karahalios, “Communicating algorithmic process in online behavioral advertising,” in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 2018, pp. 1–13.
- [64] M. B. Musa and R. Nithyanand, “Atom: ad-network tomography,” *Proceedings on Privacy Enhancing Technologies*, vol. 4, pp. 295–313, 2022.
- [65] B. Liu, A. Sheth, U. Weinsberg, J. Chandrashekar, and R. Govindan, “AdReveal: Improving transparency into online targeted advertising,” in *Proceedings of the Twelfth ACM Workshop on Hot Topics in Networks*, 2013, pp. 1–7.
- [66] P. Vallina *et al.*, “Advanced Methods to Audit Online Web Services,” Ph.D. dissertation, Universidad Carlos III de Madrid, Spain, 2023.
- [67] G. Storey, D. Reisman, J. Mayer, and A. Narayanan, “The future of ad blocking: An analytical framework and new techniques,” *arXiv preprint arXiv:1705.08568*, 2017.
- [68] X. Song, Y. Zhu, X. Zeng, and X. Chen, “Hierarchical contaminated web page classification based on meta tag denoising disposal,” *Security and Communication Networks*, vol. 2021, 2021.
- [69] Q. Zhao, W. Yang, and R. Hua, “Design and research of composite web page classification network based on deep learning,” in *IEEE 31st International Conference on Tools with Artificial Intelligence (ICTAI)*, 2019, pp. 1531–1535.
- [70] A. Gupta and R. Bhatia, “Ensemble approach for web page classification,” *Multimedia Tools and Applications*, vol. 80, no. 16, pp. 25 219–25 240, 2021.
- [71] F. Aydos, A. M. Özbayoğlu, Y. Şirin, and M. F. Demirci, “Web page classification with Google Image Search results,” *arXiv preprint arXiv:2006.00226*, 2020.
- [72] D. López-Sánchez, A. G. Arrieta, and J. M. Corchado, “Visual content-based web page categorization with deep transfer learning and metric learning,” *Neurocomputing*, vol. 338, pp. 418–431, 2019.
- [73] “Infographic: How Many Websites Are There?” <https://www.statista.com/chart/19058/number-of-websites-online>, 2023, [Accessed 28 Feb. 2023].
- [74] “Member list,” <https://www.kidsafeseal.com/certifiedproducts.html>, 2023, [Accessed 28 Feb. 2023].
- [75] “Best Websites for Kids,” <https://www.commonensemedia.org/website-lists>, 2023, [Accessed 28 Feb. 2023].
- [76] “The list of kids’ websites,” <https://kids.kaspersky.com/kids-website-list>, 2023, [Accessed 28 Feb. 2023].
- [77] “Contributors,” <https://support.virusotal.com/hc/en-us/articles/115002146809-Contributors>, 2023, [Accessed 28 Feb. 2023].
- [78] “Adding Rank Magnitude to the CrUX Report in BigQuery,” <https://developers.google.com/web/updates/2021/03/crux-rank-magnitude>, 2021, [Accessed 28 Feb. 2023].
- [79] “Fast and reliable end-to-end testing for modern web apps | Playwright,” <https://playwright.dev>, 2023, [Accessed 28 Feb. 2023].
- [80] N. Reimers and I. Gurevych, “Sentence Embeddings using Siamese BERT-Networks,” *arXiv preprint arXiv:1908.10084*, 2019.
- [81] “sentence-transformers/all-mpnet-base-v2 · Hugging Face,” <https://huggingface.co/sentence-transformers/all-mpnet-base-v2>, 2023, [Accessed 28 Feb. 2023].
- [82] K. Song, X. Tan, T. Qin, J. Lu, and T.-Y. Liu, “Mpnet: Masked and permuted pre-training for language understanding,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 16 857–16 867, 2020.
- [83] N. Reimers and I. Gurevych, “Making monolingual sentence embeddings multilingual using knowledge distillation,” *arXiv preprint arXiv:2004.09813*, 2020.
- [84] Hugging Face, “trainer · Hugging Face,” https://huggingface.co/docs/transformers/main_classes/trainer, 2023, [Accessed 28 Feb. 2023].
- [85] —, “AutoModels,” https://huggingface.co/transformers/v3.0.2/model_doc/auto.html, 2023, [Accessed 28 Feb. 2023].
- [86] Ray Project (Anyscale), “Hyperparameter Search with Transformers and Ray Tune,” <https://huggingface.co/blog/ray-tune>, 2023, [Accessed 28 Feb. 2023].
- [87] The Ray Team, “Tune: Scalable Hyperparameter Tuning — Ray 2.2.0,” <https://docs.ray.io/en/latest/tune/>, 2023, [Accessed 28 Feb. 2023].
- [88] M. Jaderberg, V. Dalibard, S. Osindero, W. M. Czarnecki, J. Donahue, A. Razavi, O. Vinyals, T. Green, I. Dunning, K. Simonyan *et al.*, “Population based training of neural networks,” *arXiv preprint arXiv:1711.09846*, 2017.
- [89] A. Stolerman, R. Overdorf, S. Afroz, and R. Greenstadt, “Classify, but verify: Breaking the closed-world assumption in stylometric authorship attribution,” in *IFIP Working Group*, vol. 11, 2013, p. 64.
- [90] M. Juarez, S. Afroz, G. Acar, C. Diaz, and R. Greenstadt, “A critical evaluation of website fingerprinting attacks,” in *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*, 2014, pp. 263–274.
- [91] scikit-learn developers, “Scikit-learn: Machine learning in Python,” https://scikit-learn.org/stable/modules/generated/sklearn.metrics.fbeta_score.html, 2011–, version 1.0.2 [Accessed 28 Feb. 2023].
- [92] “DNS0 Kids - A childproof version of the Internet.” <https://www.dns0.eu/kids>, 2023, [Accessed 4 May 2023].

- [93] “Tracker Radar Collector,” <https://github.com/duckduckgo/tracker-radar-collector>, 2023, [Accessed 28 Feb. 2023].
- [94] puppeteer, “Puppeteer,” <https://github.com/puppeteer/puppeteer>, 2023, [Accessed 28 Feb. 2023].
- [95] DuckDuckGo, “notABot.js,” <https://github.com/duckduckgo/tracker-radar-collector/blob/main/helpers/notABot.js>, 2023, [Accessed 3 Aug. 2023].
- [96] “Early browser API accesses and function calls are missed,” <https://github.com/duckduckgo/tracker-radar-collector/issues/77>, 2023, [Accessed 2 Aug. 2023].
- [97] “Browserleaks - Check your browser for privacy leaks,” <https://browserleaks.com>, 2023, [Accessed 3 Aug. 2023].
- [98] “@gorhill/ubo-core,” <https://www.npmjs.com/package/@gorhill/ubo-core>, 2023, [Accessed 28 Feb. 2023].
- [99] R. Hill and N. Rolls, “uBlock Origin,” <https://ublockorigin.com/>, [Accessed 28 Feb. 2023].
- [100] —, “uBlock,” <https://github.com/gorhill/uBlock/blob/491bc87e94a503a17fd11cdee35cf1b6fea24be/platform/mv3/make-rulesets.js#L1285-L1296>, 2023, [Accessed 28 Feb. 2023].
- [101] “entity_map.json - DuckDuckGo Tracker Radar,” https://github.com/duckduckgo/tracker-radar/blob/main/build-data/generated/entity_map.json, [Accessed 28 Feb. 2023].
- [102] T. Urban, M. Degeling, T. Holz, and N. Pohlmann, “Beyond the front page: Measuring third party dynamics in the field,” in *Proceedings of The Web Conference 2020*, 2020, pp. 1275–1286.
- [103] W. Aqeel, B. Chandrasekaran, A. Feldmann, and B. M. Maggs, “On Landing and Internal Web Pages: The Strange Case of Jekyll and Hyde in Web Performance Measurement,” in *Proceedings of the ACM Internet Measurement Conference*, 2020, pp. 680–695.
- [104] A. Rasaii, S. Singh, D. Gosain, and O. Gasser, “Exploring the cookieverse: A multi-perspective analysis of web cookies,” in *International Conference on Passive and Active Network Measurement*. Springer, 2023, pp. 623–651.
- [105] DuckDuckGo, “autoconsent,” <https://github.com/duckduckgo/autoconsent>, 2023, [Accessed 28 Feb. 2023].
- [106] “Web Tracking Protections — DuckDuckGo Help Pages,” <https://help.duckduckgo.com/duckduckgo-help-pages/privacy/web-tracking-protections/#cookie-pop-up-management>, 2023, [Accessed 3 May 2023].
- [107] R. Bagge, C. Matte, ric Daspet, K. Emanuel, S. Macbeth, and S. Roeland, “Consent-O-Matic,” <https://github.com/cavi-au/Consent-O-Matic/>, 2019, [Accessed 28 Feb. 2023].
- [108] M. Nouwens, I. Liccardi, M. Veale, D. Karger, and L. Kagal, “Dark Patterns after the GDPR: Scraping Consent Pop-Ups and Demonstrating Their Influence,” in *CHI Conference on Human Factors in Computing Systems*, 2020, pp. 1–13.
- [109] The EasyList authors, <https://easylist.to/easylist/easylist.txt>, 2023, [Accessed 28 Feb. 2023].
- [110] “Transparency and ad disclosures - Google Ads Help,” <https://support.google.com/google-ads/answer/9729263?hl=en>, 2023, [Accessed 28 Feb. 2023].
- [111] L. Gak, S. Olojo, and N. Salehi, “The Distressing Ads That Persist: Uncovering The Harms of Targeted Weight-Loss Ads Among Users with Histories of Disordered Eating,” *Proc. ACM Hum.-Comput. Interact.*, vol. 6, no. CSCW2, Nov 2022.
- [112] A. J. Campbell, “Rethinking children’s advertising policies for the digital age,” *Loy. Consumer L. Rev.*, vol. 29, p. 1, 2016.
- [113] “Age-restricted ads online,” <https://www.asa.org.uk/static/44dc1935-0766-4378-91171e6954ae560a/Age-restricted-ads-online-targeting-guidance.pdf>, 2023, [Accessed 28 Feb. 2023].
- [114] “Ads & Made for Kids content,” <https://support.google.com/adspolicy/answer/9683742?hl=en>, 2023, [Accessed 1 Mar. 2023].
- [115] “Restricted Content, Products, and Services,” <https://help.taboola.com/hc/en-us/articles/115000220793-Restricted-Content-Products-and-Services>, 2023, [Accessed 1 Mar. 2023].
- [116] G. Inc., “Google Cloud Vision API’s SafeSearch Detection,” <https://cloud.google.com/vision/docs/detecting-safe-search>, 2021.
- [117] Microsoft, “Content Moderator - Image Moderation,” <https://docs.microsoft.com/en-us/azure/cognitive-services/content-moderator/image-moderation-api>, 2021.
- [118] M. Grootendorst, “Bertopic: Neural topic modeling with a class-based tf-idf procedure,” *arXiv preprint arXiv:2203.05794*, 2022.
- [119] “IngoSearch,” <https://ingosearch.com>, 2023, [Accessed 4 May 2023].
- [120] “BetterMe,” <https://betterme.world>, 2023, [Accessed 4 May 2023].
- [121] “An unhealthy diet of targeted ads: an investigation into how the diet industry exploits our data,” <https://privacyinternational.org/long-read/4603/unhealthy-diet-targeted-ads-investigation-how-diet-industry-exploits-our-data>, 2021, [Accessed 5 May 2023].
- [122] “Eis.de Is Now Germany’s Biggest Adult Retailer,” <https://www.venus-adult-news.com/en/web-tech/e-commerce/eis-de-is-now-germanys-biggest-adult-retailer>, 2023, [Accessed 3 Aug. 2023].
- [123] S. Bell, “How URL Tracking Systems are Abused for Phishing,” *Security Boulevard*, 2020, <https://securityboulevard.com/2020/10/how-url-tracking-systems-are-abused-for-phishing>.
- [124] “Pixel for Marketing API - Meta Pixel - Documentation - Meta for Developers,” <https://developers.facebook.com/docs/meta-pixel/implementation/marketing-api>, Aug. 2023, [Accessed 2 Aug. 2023].
- [125] “LinkedIn Insight Tag | LinkedIn Marketing Solutions,” <https://business.linkedin.com/marketing-solutions/insight-tag>, Aug. 2023, [Accessed 3 Aug. 2023].
- [126] “Advanced Store,” <https://www.advanced-store.com/en/>, [Accessed: February 24, 2023].
- [127] “Solutions,” <https://throttle.io/solutions>, 2021, [Accessed 31 Jul. 2023].
- [128] “Stripe | Payment Processing Platform for the Internet,” <https://stripe.com/en-nl>, 2023, [Accessed 28 Feb. 2023].
- [129] “Join the Smart Affiliate Marketing Network | WEBGAINS,” <https://www.webgains.com/public/en>, 2023, [Accessed 28 Feb. 2023].
- [130] European Data Protection Board, “Binding Decision 4/2022 on the dispute submitted by the Irish SA on Meta Platforms Ireland Limited and its Instagram service (Art. 65 GDPR),” https://edpb.europa.eu/our-work-tools/our-documents/binding-decision-board-art-65/binding-decision-42022-dispute-submitted_en, 2022, [Accessed 28 Feb. 2023].
- [131] ePrivacy Directive, “Directive 2002/58/EC of the European Parliament and of the Council,” <http://data.europa.eu/eli/dir/2002/58/2009-12-19>, 2009, [Accessed 28 Feb. 2023].
- [132] I. Milkaite and E. Lievens, “Status quo regarding the child’s article 8 GDPR age of consent for data processing across the EU,” *BIK PORTAL*, no. 20/12/2019, 2019.
- [133] European Data Protection Board, “Guidelines 05/2020 on consent under Regulation 2016/679,” https://edpb.europa.eu/sites/default/files/files/file1/edpb_guidelines_202005_consent_en.pdf, 2020, [Accessed 28 Feb. 2023].
- [134] “Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market for Digital Services and amending Directive 2000/31/EC (Digital Services Act),” <http://data.europa.eu/eli/reg/2022/2065/oj>, 2022, [Accessed 28 Feb. 2023].
- [135] European Commission, “Digital Services Act: Commission designates first set of Very Large Online Platforms and Search Engines,” https://ec.europa.eu/commission/presscorner/detail/en/IP_23_2413, 2023.

- [136] FTC, “Complying with COPPA: Frequently Asked Questions,” <https://www.ftc.gov/node/60394>, 2020, [Accessed 28 Feb. 2023].
- [137] —, “Fortnite Video Game Maker Epic Games to Pay More Than Half a Billion Dollars over FTC Allegations of Privacy Violations and Unwanted Charges,” <https://www.ftc.gov/node/80135>, 2022, [Accessed 28 Feb. 2023].
- [138] “Digital Advertising - Worldwide | Statista Market Forecast,” <https://www.statista.com/outlook/dmo/digital-advertising/worldwide>, 2023, [Accessed 28 Feb. 2023].
- [139] “5Rights,” <https://5rightsfoundation.com>, Nov. 2023, [Accessed 29 Nov. 2023].
- [140] D. Zeber, S. Bird, C. Oliveira, W. Rudametkin, I. Segall, F. Wollsen, and M. Lopatka, “The representativeness of automated web crawls as a surrogate for human browsing,” in *Proceedings of The Web Conference 2020*, 2020, pp. 167–178.
- [141] J. Jueckstock, S. Sarker, P. Snyder, A. Beggs, P. Papadopoulos, M. Varvello, B. Livshits, and A. Kapravelos, “Towards realistic and reproducible web crawl measurements,” in *Proceedings of the Web Conference 2021*, 2021, pp. 80–91.
- [142] M. L. McHugh, “Interrater reliability: the kappa statistic,” *Biochemia Medica*, vol. 22, no. 3, pp. 276–282, 2012.

Appendix A. Criteria for labeling child-directed websites

To identify a child-directed website, we manually visit and review its design, content, and policies, including necessary translations to English. A site is labeled as child-directed if any of the following conditions are met:

- Does the website include content, activities, or games that can be used by children?
- Does the website promote products (e.g., apps, sites, books, videos, workshops, animations, etc.) designed for and usable by children online?
- Does the website include content or promote products whose end users are children, but children’s parents must first subscribe or register?

A site is not child-directed if one of the following is true:

- The website redirects to another page that is not child-directed.
- The website features children-related products intended for adult use, such as parents or teachers.
- The website is generally appealing to adults (e.g., news or academic websites).

A.1. Manual Verification of the Classifier Output

Two researchers manually labeled a total of 2,500 websites detected as child-directed by the classifier. Initially, both researchers jointly labeled a 50-website sample, reaching agreement on 45 of these decisions (Cohens Kappa=0.79 [142]). The remaining websites were divided into two equal batches and labeled individually by the two researchers. This process took approximately one person-week to complete. We followed the criteria for identifying child-directed websites (Appendix A) and considered four potential labels for each website: *child-directed*, *child-related*, *non-children* and *unknown*. The majority of the

websites (64%) were labeled as child-directed, , while 23% were identified as child-related, indicating they are relevant to children but primarily intended for parents or teachers. About 4.5% of the websites were inaccurately classified as non-children’s sites, such as entertainment bands, academic forums, and movie streaming services, mostly due to vague or brief titles and descriptions. In some cases, discerning whether websites were targeted at children, parents, or teachers was challenging, leading to 5.5% being labeled as ‘unknown’ due to uncertainty about their target audience. Of the misclassified websites, we found that four were adult entertainment websites (0.16%) that had very short metadata fields mentioning words such as “teens”, “cartoon”, “animations” which likely caused the misclassification.

A.2. Ad Transparency Statements

The following ad transparency statements are used to classify advertisements as targeted or non-targeted. Note that targeted categories also include retargeting and behavioral ads. The statements are compiled from Google’s and Criteo’s ad disclosure interfaces, reached via the AdChoices icon. When searching for the statements, we use exact, case-insensitive search.

Targeted:

- Google’s estimation of your interests
- Websites you’ve visited
- Your similarity to groups of people the advertiser is trying to reach
- Your activity on Google on this device
- According to your activity on this device
- You have enabled ad personalization
- Information collected by the publisher. The publisher partners with Google to show ads
- Google’s estimation of the languages you know, based on your activity on this device
- Your visit to the advertiser’s website or app
- The advertiser’s interest in reaching new customers who haven’t bought something from them before

Non-Targeted:

- Ad personalization is turned off
- You have turned off ad personalization
- Ad personalization is off
- The time of day or your general location
- Google shows ads based on general factors like the time of day and the info on a page, our policies, and your ad personalization settings
- The information on the website you were viewing
- General factors about the placement of the ad

A.3. Detecting failed or errored visits

We classified a visit as failed under the following conditions: if the first request elicited a 4XX or 5XX error; if the size of the first non-3XX response (root document) was less than 512 bytes, following Le Pochat et al. [26]; or if there was no successful (200 OK) response.

Appendix B. Meta-Review

The following meta-review was prepared by the program committee for the 2024 IEEE Symposium on Security and Privacy (S&P) as part of the review process as detailed in the call for papers.

B.1. Summary

This paper examines the presence of both tracking and inappropriate ads on websites directed at children. The authors built a pipeline to identify children's websites and collect their advertisements, which prior work achieved manually. After collecting these advertisements, the paper also demonstrates that in addition to a high prevalence of trackers which should require informed parental consent, many children's sites also display inappropriate ads.

B.2. Scientific Contributions

- Independent Confirmation of Important Results with Limited Prior Research
- Provides a New Data Set For Public Use
- Creates a New Tool to Enable Future Science
- Provides a Valuable Step Forward in an Established Field
- Other

B.3. Reasons for Acceptance

- 1) This paper makes novel progress on studying tracking and targeted advertising as it relates to children, particularly in its dramatic increase in scope compared to prior work.
- 2) The authors provide a new dataset of child-directed websites that is much larger than prior work, along with the other data and code used for analysis in regards to better measuring tracking and targeted ads.

B.4. Noteworthy Concerns

- 1) "Improper" as an ad criterion is not clearly defined. Although the authors reference prior work on this topic, further discussion of how and why they picked the specific four criteria used in the paper would be useful for contextualizing their results.
- 2) One reviewer noted that the recall of child-directed websites was rather low at 70%, leading to potential skew in the analysis.

Appendix C. Response to the Meta-Review

- 1) Our analysis of improper ads is exploratory, and it serves as an initial step toward a more extensive investigation into the problematic advertising

practices. We focus on four types of improper ads (weight loss, mental health, dating services, and racy content), because prior research [111], [112] and regulatory guidelines [113] have argued that these can be especially harmful to children. Moreover, these categories are generally restricted by ad networks own moderation policies [114], [115].

While our study focuses on these four categories, our versatile ML improper ad detection pipeline can be easily adapted to detect different kinds of improper ads, for instance, by simply changing the search queries (Figure 4).

- 2) The lower recall rate is the result of employing the "Classify-Verify" method in our classifier. This approach enhances precision by applying a threshold, which leads to lower recall but effectively reduces false positives and the manual labeling workload (§3.2). In particular, following Juarez et al. [90], we choose the threshold that maximizes $F_{\beta=0.5}$, which gives more weight to precision to reduce false positives. Researchers who adopt our methods can easily tweak the threshold to strike a different balance based on their classification needs and manual verification budgets.

We acknowledge that selective (abstaining) classification may have a disparate effect on websites that are classified with low confidence. Other factors including the website language, title and descriptions may also skew the results, as we note in the Limitations. We acknowledge these limitations and do not claim that our list of child-directed websites is representative of the child-directed websites at large.