109

110

111

112

113

114

115

116

59



Anonymous Author(s)

Abstract

1

2

3

5

8

9

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

26

27

28

29

30

31

32

33

34

35

36

37

38

39

40

41

42

43

44

45

46

47

48

49

50

51

52

53

54

55

Browser fingerprinting is a pervasive online tracking technique increasingly used for profiling and targeted advertising. Existing research on fingerprinting prevalence relies heavily on automated web crawls, which inherently struggle to replicate the nuances of human-computer interaction. This raises concerns about the accuracy of current understandings of real-world fingerprinting deployments. To that end, this paper presents a user study involving 30 participants over a 10-week period, capturing telemetry data from real browsing sessions across 3,000 top-ranked websites.

Our findings reveal that automated crawls miss nearly half (47.8%) of the fingerprinting websites encountered by real users. This discrepancy mainly stems from crawlers' inability to access authentication-protected pages, circumvent bot detection mechanisms, and trigger fingerprinting scripts activated by specific user interactions. We also identify potential new fingerprinting vectors present in real user data but absent from automated crawls. Finally, we evaluate the effectiveness of federated learning for training browser fingerprinting detection models on real user data, demonstrating superior performance to models trained solely on automated crawl data.

CCS Concepts

• Security and privacy → Privacy-preserving protocols; Privacy protections; • Networks → Web protocol security.

Keywords

Browser Fingerprinting; Differential Privacy; Federated Learning

ACM Reference Format:

Anonymous Author(s). 2024. Beyond the Crawl: Unmasking Browser Fingerprinting in Real User Interactions. In *Submission to The Web Conference (TheWebConf'25), April 28–May 2, 2025, Sydney, Australia.* ACM, New York, NY, USA, 12 pages.

1 Introduction

Browser fingerprinting is an invasive online tracking technique widely considered a significant threat to users' privacy [9, 26, 55]. Generally, it involves trackers using client-side device information (e.g., hardware specifications or browser configurations) to derive unique device identifiers and track users across multiple visits and websites. Unlike third-party cookies, browser fingerprinting is *stateless* and thus less visible to users, who have limited control and few tools to defend against it.

Worse yet, browser fingerprinting is highly intrusive, as the identifiers remain stable over long periods of time [48] and can

Please use nonacm option or ACM Engage class to enable CC licenses This work is licensed under a Creative Commons Attribution 4.0 International License. TheWebConf'25, April 28–May 2, 2025, Sydney, Australia

© 2024 Copyright held by the owner/author(s).

56 57 58 be effective even when using incognito mode [4]. Countermeasures can be tricky to deploy after detection, as they often lead to significant website breakage and affect user experience [17, 31], not to mention that the unique identifier may have already been revealed at that point. With the increasing prevalence of fingerprinting [31] and the phasing out of third-party cookies by major browsers, countering invasive browser fingerprinting has become crucial to protecting user privacy on the modern Web.

Detecting fingerprinting attempts, often implemented through JavaScript scripts, is a prerequisite for effective mitigation. Initially, these scripts were mainly identified through simple heuristics [1, 27] and manually curated blocklists [20, 23, 25]; however, these are hard to maintain and often narrowly defined to reduce false positives, thus missing many fingerprinting scripts in practice. Consequently, machine learning (ML) based detectors have become increasingly popular [18, 29, 31] and can detect significantly more fingerprinting scripts with comparable false positive rates to heuristics. Overall, these techniques generally rely on centralized, automated crawlers instructed to visit a large number of websites to collect scripts for analysis and fingerprinting detection.

One major limitation of automated crawlers is their limited ability to faithfully replicate genuine user behaviors and interactions. Although recent work has made progress toward emulating some degree of user interactions (such as accepting cookie banners [34, 51] and emulating user interests in browsing patterns [34]), many key actions remain firmly outside of the reach of automated crawlers. For instance, these often fail to reliably solve CAPTCHAs, evade bot detectors, login, access websites behind paywalls, etc. This leads to incomplete website coverage and potentially missed fingerprinting scripts. For example, Iqbal et al. [31]'s automated crawl failed to visit 11.9% of top-ranked websites.

To address this limitation, prior work has investigated the feasibility of training ML models on real-user browsing sessions. Evidently, this needs to be done in a privacy-preserving way to avoid exposing users' browsing histories, etc. Annamalai et al. [8] recently proposed FP-Fed, a system relying on differentially private federated learning (DP-FL) to collaboratively train ML models on the combined browsing sessions of many users while providing strong formal privacy guarantees. Although FP-Fed can train models that achieve good performance even at moderate levels of privacy, the authors only tested it on a dataset derived from an automated crawl, with unclear implications for its performance in the real world.

Technical Roadmap. In this paper, we investigate the prevalence and distribution of browser fingerprinting in real-user browsing sessions, as opposed to automated crawls used in prior work [8, 27, 31]. Specifically, we collect real-world data from a 10-week study (June-August 2024) involving 30 participants who browsed the Web with their own devices and reported telemetry via a Chrome extension. This enabled us to collect browser fingerprinting signals from real browsing sessions across 3,000 top-ranked websites.

118

119

120

121

123

134

135

136

137

138

139

140

141

142

143

144

145

146

147

148

149

150

151

152

153

154

155

156

157

158

159

160

161

162

163

164

165

166

167

174

This study allows us to shed light on two distinct aspects of fingerprinting. First, we compare the prevalence and distribution of browser fingerprinting in a real-world dataset to an equivalent automated crawl of the same 3,000 websites. Second, we conduct a comparative analysis of the ML performance of browser fingerprinting detectors trained on the automated crawl and the real-world dataset in a distributed and privacy-preserving way.

Main Findings. Overall, our study and experimental analysis yield
 three interesting findings:

- Surprisingly, almost half(47.8%) of fingerprinting websites identified from real-user browsing sessions are missed by automated crawls. This is due to three main reasons: (1) authentication pages that are typically not crawled by automated crawlers, (2) bot detection scripts that track real user interaction before browser fingerprinting is triggered, and (3) cookie banners that require user consent before performing fingerprinting.
 We have a state of the s
 - We discover potential new fingerprinting vectors from the realuser browsing sessions that were not previously found in the automated crawls.
 - We show that ML models trained privately on real-user browsing sessions can achieve comparable or even better performance than non-private models trained on automated crawls alone while providing privacy and learning the behaviors of many more fingerprinting scripts. Specifically, in our experiments, the former achieves an Area Under the Precision-Recall Curve (AUPRC) of 0.982 at a privacy level of $\varepsilon = 5$, compared to an AUPRC of 0.971 for the latter.

Overall, our work paves the way for large-scale deployment of more effective, dynamic, and robust browser fingerprinting detection relying on a scalable, distributed, and privacy-preserving infrastructure geared to be readily integrated into modern browsers.

2 Background & Related Work

In this section, we review browser fingerprinting and differentially private federated learning, along with relevant prior work.

2.1 Browser Fingerprinting

Browser fingerprinting is a tracking technique usually deployed through Javascript running on a user's browser to build a unique user identifier. This user identifier typically consists of high-entropy device information (e.g., screen size, GPU model) used to produce highly unique and stable user identifiers. Unlike third-party cookies, browser fingerprinting is *stateless*, i.e., no data is stored on the user's device and, therefore, cannot be easily detected or mitigated. Although browser fingerprinting has been known to be used for legitimate purposes, e.g., Web authentication [5, 33, 51] or fraud detection [30, 35], it is often used for online tracking and to serve targeted ads [34]. As a result, browser fingerprinting is widely considered a significant threat to user privacy [55], thus prompting many browser vendors to deploy countermeasures [9, 12, 26].

Although there are well-documented cases of fingerprinting in the wild, we are not aware of a widely accepted formal definition of fingerprinting [31]. Mayer [36] was the first to observe that the uniqueness and customization of browsing environments (which they called "quirkiness") can be abused to identify users. The largescale Panopticlick experiment later conducted by the late Peter Eckersley [24] showed that most browsing sessions (83.6%) have unique fingerprints. While these results are concerning on their own, the fingerprinting surface considered by this early work was relatively limited to information collected from simple Javascript APIs like the Screen and Navigator APIs and HTTP headers. More recently, as more features and APIs were added to the Javascript specification, fingerprinting has begun to include Canvas [40], Audio [27], WebRTC [27], WebGL [13], Battery Status [45], Mobile Sensor [11], and even the Web Bluetooth APIs [10]. With such a large fingerprinting surface, it is often difficult to pinpoint the *intent* behind the use of these various APIs in arbitrary Javascript scripts, thus making it challenging to specify a single comprehensive definition of fingerprinting.

Nevertheless, we follow prior work by Iqbal et al. [31] and take a conservative approach to defining fingerprinting based on a well-known high precision (low false positive rate) heuristic. Specifically, we exclude the simple curation of properties from the Navigator and Screen APIs and focus on the four most obvious forms of fingerprinting (Canvas, Canvas Font, Audio, and WebRTC). This definition not only minimizes the probability of flagging non-fingerprinting scripts but is also helpful in training ML models that generalize well to other forms of fingerprinting as well [8].

Next, we briefly describe the four main types of fingerprinting identified by the heuristics [27] and their associated criteria.

Canvas. Scripts exploit differences in how fonts are rendered across devices. Criteria:

- text is written to the canvas element using the fillText or strokeText method;
- (2) style is applied with fillStyle or strokeStyle method;
- (3) toDataURL is called to extract the image from the canvas; and
- (4) save, restore or addEventListener methods are not called.

Canvas Font. Scripts access the list of fonts installed on a device. Criteria:

- font property of canvas element is set to more than 20 different fonts; and
- (2) measureText is called more than 20 times.

WebRTC. Scripts rely on the uniqueness of peers in the WebRTC protocol [41]. Criteria:

- createDataChannel or createOffer method is called on a WebRTC peer connection; and
- (2) onicecandidate or localDescription method is called.

AudioContext. Scripts exploit differences in the different hardware processing audio. Criteria:

 either createOscillator, createDynamicsCompressor, destination, startRendering, or oncomplete are called.

2.2 Detecting Browser Fingerprinting

Early research on browser fingerprinting detection has mainly relied on manually curated blocklists, e.g., EasyPrivacy [23], Privacy Badger [25], and Disconnect [20]. However, as fingerprinting vendors and scripts constantly evolve, it can be challenging to maintain these lists continuously.

Heuristics-based detection. Manual analysis of fingerprinting scripts by Acar et al. [2] and Roesner et al. [49] paved the way

175

176

177

178

179

180

181

182

183

184

185

186

187

188

189

190

191

192

193

194

195

196

197

198

199

200

201

202

203

204

205

206

207

208

209

210

211

212

213

214

215

216

217

218

219

220

221

222

223

224

225

226

227

228

229

230

231

for the first heuristic that detected canvas fingerprinting automat-233 ically [1]. More precisely, Acar et al. [2] monitored and analyzed 234 235 the arguments and return values of the fillText, strokeText, and toDataURL methods exposed by the Canvas API. Three addi-236 tional types of fingerprinting (Canvas Font, WebRTC, and Audio) 237 were later added by Englehardt and Narayanan [27]. This combined 238 heuristic is now widely used as a prominent indicator of finger-239 printing [8, 18, 31, 51] as it is known to produce very few false 240 241 positives, which is an important consideration to prevent falsely 242 flagging fingerprinting scripts. However, as noted by Iqbal et al. [31], it might miss many fingerprinting scripts in practice since it 243 is defined very narrowly to achieve high precision. Furthermore, 244 keeping the heuristic up-to-date with the latest Javascript APIs and 245 fingerprinting vectors can be challenging. 246

247 ML-based detection. Methods based on machine learning solve 248 the problem of manually maintaining blocklists and heuristics by 249 learning fingerprinting behaviors in the wild. Ikram et al. [29] 250 trained a one-class SVM on static features directly extracted from 251 scripts' source code. However, Iqbal et al. [31] observed that code is 252 often obfuscated, making it difficult to reliably learn fingerprinting 253 behaviors from static features alone; therefore, they additionally 254 trained a Decision Tree classifier on dynamic features extracted 255 from the execution trace of a given script. By monitoring and ana-256 lyzing the number of times a script calls each Javascript API, along 257 with the associated arguments and return values, Iqbal et al. [31] 258 trained a robust browser fingerprinting detector that achieves both 259 high precision and high recall.

260

261

262

263

264

265

266

267

268

269

270

271

272

273

274

275

276

277

278

279

280

281

282

283

284

285

286

287

288

289

290

Recently, Annamalai et al. [8] proposed going beyond *centralized* models that rely on a large automated crawl, which cannot replicate human interaction and, therefore, might miss fingerprinting scripts in the wild. They proposed FPFed, a system geared to train a browser fingerprint detection model collaboratively on real users' browsing sessions while preserving privacy using the federated learning paradigm [37] (see below). However, Annamalai et al. [8] did not actually evaluate their system on real-world user behavior but opted to test their system by simulating real-world users on an automated crawl. Our work aims to fill this research gap by collecting browsing data from real users and evaluating the FP-Fed system on a realworld dataset instead.

2.3 Differentially Private Federated Learning

The standard way to train models collaboratively on multiple users' data while providing formal privacy guarantees is through differentially private federated learning (DP-FL). More precisely, DP-FL combines Federated Learning (FL) and Differential Privacy (DP).

Federated Learning. Typically, centralized ML models are trained on datasets stored at a single entity. However, in many scenarios, gathering data from multiple users may not be possible due to privacy, security, and/or efficiency concerns. To this end, Federated Learning (FL) [37] introduces a distributed learning approach allowing users to train an ML model collaboratively without disclosing their (potentially sensitive) training data.

In FL, users train local models on their individual training data, sharing only their model updates (and not the raw training data) with a server. The server aggregates the model updates to build a global model, which is then propagated back to the users. This process then repeats until the global model converges. Although FL ensures that the server never sees the raw training data, prior work has shown that the model updates can still leak sensitive information about users' training data [39]. Therefore, recent work has focused on providing formal privacy guarantees.

Differential Privacy (DP). DP is the standard framework for defining algorithms that provide theoretical upper bounds on the loss of privacy incurred by data subjects due to the output of an algorithm [22].

Definition 2.1 (Differential Privacy (DP)). A randomized mechanism $\mathcal{M} : \mathcal{D} \to \mathcal{R}$ is (ε, δ) -differentially private if for any two neighboring datasets $D, D' \in \mathcal{D}$ and $S \subseteq \mathcal{R}$

$$\Pr[\mathcal{M}(D) \in S] \le e^{\varepsilon} \Pr[\mathcal{M}(D') \in S] + \delta$$

The privacy guarantees provided by a differentially private algorithm are characterized by the parameters ε and δ . The privacy parameter, ε , is often referred to as the privacy budget and ranges from 0 to ∞ , with lower values denoting better privacy. Whereas δ quantifies the probability that the mechanism fails to deliver any guarantees; typically, δ is fixed to some asymptotically small number (e.g., 10^{-5}).

DP-FL. DP can be combined with FL in many ways to provide strong, formal privacy guarantees in the FL setting. One way is by adding statistical noise when the server aggregates the model updates, aka Central DP (CDP). Under this regime, the model is trusted with the aggregated model updates but not the raw sensitive data. CDP guarantees that it is impossible (up to the privacy parameter ε) to infer whether or not data from a user was used to train the global model based on the aggregated and noised model updates. Examples of CDP instantiations for DP-FL include next-word prediction [38], medical image analysis [3], and network analysis [42]. Other approaches use Local DP [52, 54] or Distributed DP [32] to minimize trust assumptions in the server.

In this work, we consider the FP-Fed [8] setting, which uses CDP, as it provides a good trade-off between utility, convergence speed, and the amount of noise required for the desired level of privacy. (A detailed overview of the FP-Fed system can be found in Appendix A.)

3 Methodology

Next, we discuss our methodology to collect browser fingerprinting scripts from real users' browsing sessions. We describe the websites we collect data from and how we recruit participants, collect website telemetry, and detect fingerprinting. We also discuss relevant ethical considerations with respect to collecting and analyzing the data.

3.1 Websites of Interest

Ideally, to capture real user interactions, we would want to collect and analyze telemetry data from all websites visited by participants. However, browsing histories are highly sensitive, and collecting them might violate participants' privacy, discourage them from contributing to the study, or even introduce biases in the data collection. (Note that while FL-based detection would ensure that raw telemetries are not disclosed, we do need them for our experiments when measuring accuracy and finding discrepancies).

350

351

352

353

354

355

356

357

358

359

360

361

362

363

364

365

366

367

368

369

370

371

372

373

374

375

376

377

378

379

380

381

382

383

384

385

386

387

388

389

390

391

392

393

394

395

396

397

398

406

As a result, we opt to collect telemetry from a selected set of websites-more specifically, the top-ranked websites according to the Chrome User Experience Report (CrUX) from April 2024 [14]. Overall, we collect data from 3K websites sampled from the top 5K ranked websites in the CrUX report: specifically, we take the top 1K ranked websites and a random 2K sample of top 1K to 5K ranked websites. To further limit the amount of sensitive information collected, we use Cloudflare's Domain Intelligence API [15] to detect and filter out websites containing adult or potentially harmful content. Specifically, we exclude websites from the following categories (as defined by the Domain Intelligence API): 1) *Adult Themes*, 2) *Gambling*, 3) *Questionable Content*, 4) *Security Threats*, 5) *Violence*, and 6) *Security Risks*. Out of the top 5K ranked websites from the CrUX ranking, we excluded 948 such websites, before sampling the 3K websites that we use in our study.

3.2 Participant Recruitment

We recruited 30 participants (aged 18 to 60) who used the Chrome browser through the Amazon Mechanical Turk (MTurk) platform [6]. The "Human Intelligence Task" (HIT), i.e., the ad used to recruit users on MTurk, can be found in Appendix B.

Each participant was instructed to download a Chrome extension we developed and given a password to authenticate on the extension. As mentioned above, for privacy purposes, each participant was provided with a pre-defined set of 100 websites to visit instead of collecting data from their "natural" browsing sessions. (Note that the extension only collects data from this list.)

For each website, the participants were instructed to visit at least ten sub-pages (including login pages), accept cookie banners, and solve CAPTCHAs when presented to simulate realistic browsing sessions. The extension collected the telemetry from the websites in the background and transmitted it to our server. Other than the website telemetry required for the purpose of fingerprinting detection, no other data (e.g., demographics, name, email) was collected. Upon successfully visiting all of the given websites, the participants were given a "Task Completion Code" by the extension, which was then used to compensate them through MTurk. In total, the participants received an average of USD 20.33 (excluding taxes and fees), which exceeds the federal minimum wage.

Although only a relatively small number of participants were recruited, we emphasize that the focus of this study is not to draw conclusions about the distribution or popularity of websites as visited by real users in their natural browsing behaviors. Specifically, the intent of this study is to investigate to what extent automated crawlers miss fingerprinting scripts due to their inability to replicate real human interactions in practice. To that end, this study could have been equivalently performed by a single participant crawling all 3K websites, but this would have made it hard to recruit a participant to complete the task successfully.

Data Quality. Overall, the participants visited a total of 14,895 unique URLs. In Figure 1, we plot the distribution of unique URLs visited by each participant. While the extension verified that each participant visited all of the websites assigned to them before providing the "Task Completion Code," we did not verify if each participant did visit at least ten sub-pages on each website as this would have been difficult to do so in general (e.g., single page applications



Figure 1: Distribution of unique URLs visited by each study participant.

may not redirect to different URLs). On average, each participant visited 506 unique URLs. This is five times the number of unique URLs visited by automated crawlers in prior work [27, 31].

3.3 Collecting Scripts and Extracting Features

Key to identifying browser fingerprinting is dynamically analyzing scripts that are loaded by the websites [31, 43]. To collect this data, as mentioned, we built a Chrome extension that monitors and records the Javascript APIs accessed by each script loaded on a website, along with the associated arguments and processed return values. Our extension then processes the collected data and sends telemetry to our server, where we then analyze it and detect fingerprinting scripts. The extension only collects data from the domains corresponding to the list of websites assigned to each participant, i.e., it does not gather data from other websites, which the participant might visit in other browsing sessions.

We built our extension based on the instrumentation developed by Iqbal et al. [31] for automated crawls with the Mozilla Firefox browser. We adapted the instrumentation to work with the Google Chrome browser instead, extending it to monitor Chrome-specific APIs as well, and injected the instrumentation using a Chrome extension that participants can easily install.

Additionally, to reduce storage and data transmission costs, our extension first pre-processes the raw data and extracts only the features necessary to detect fingerprinting. These features consist of API call counts (i.e., how many times each Javascript API is called) and "custom" features processed from the arguments and return values (e.g., length of string argument, number of elements in returned list value). This pre-processing step not only reduces the data transmission costs but also ensures that the collected data is more private, as the exact arguments and return values are not sent to the server. In summary, our extension processes the websites visited by the participant, the scripts loaded by each website, extracts features from each script's execution trace, and sends this data to our server for analysis.

3.4 Fingerprinting Detection

We follow prior work on browser fingerprinting detection [8, 27, 31] and use high-precision heuristics to label scripts as *fingerprinting*. Specifically, we use the heuristic developed by [27] and later modified by [31]. We follow this conservative approach to labeling fingerprinting scripts over machine learning classifiers [31] to ensure low false positive rates and have good confidence that the labeled scripts are, in fact, fingerprinting. The heuristics we use identify four main types of fingerprinting – namely, Canvas, Canvas Font, WebRTC, and Audio Context (see Section 2.1 – and does not

consider simple accesses to device information as fingerprinting to minimize false positives.

3.5 Ethics

465

466

467

468

469

470

471

472

473

474

475

476

477

478

479

480

481

482

483

484

499

500

501

502

503

504

505

506

507

508

509

510

511

512

522

Our study was reviewed and approved by our Institutional Review Board (details are omitted to preserve submission anonymity). As part of the process, we submitted the full documentation detailing our approach to recruiting and compensating participants, the types of data that will be collected, along with data access policies, potential ethical issues, mitigation strategies, as well as the information sheet and consent form provided to participants.

All participants were recruited and compensated entirely anonymously through the Amazon MTurk platform. Specifically, each participant received an average of USD 20.33 and was conservatively expected to take at most an hour to complete visiting all 100 websites assigned to them. Therefore, the compensation was well above minimum wages both in California [44] (USD 16/hour) and the United Kingdom [16] (GBP 11.44/hour).

In line with data minimization principles, we only collected pre-485 486 processed and extracted features from the scripts' execution traces. 487 Specifically, we did not collect the raw return values and arguments of the Javascript APIs and did not collect any other user data or 488 metadata, e.g., IP address, device, or network information. 489

As mentioned, we provided participants with a list of websites to 490 visit instead of monitoring their "natural" browsing sessions, and de-491 signed our Chrome extension to only collect data from this list. That 492 is, on all other websites, the extension does not inject the instru-493 mentation script and does not monitor or record any information 494 from these websites. By doing so, we can filter out potentially em-495 496 barrassing or dangerous websites (e.g., adult entertainment, spam) 497 and prevent accidentally collecting data from participants' visits to such websites. 498

Additionally, the exact types of data collected along with the purpose of the data collection, potential privacy implications of taking part in the study, and participants' data rights (e.g., withdrawal from data collection) were explained in layman terms through a participant information sheet which they could download. Furthermore, explicit user consent was obtained through a consent form, which reiterated their rights and privacy implications of taking part. Overall, we abided by a strict code of ethics (i.e., RESPECT [28]). In line with our institution's data retention policies, we will delete all data within ten years of the publication of our results.

4 Results

In this section, we present the results of our study. We begin by 513 comparing the prevalence of browser fingerprinting found in real 514 user browsing sessions with that observed in an automated crawl of 515 516 the same 3K top-ranked websites. Next, we shed light on why some fingerprinting websites are missed by automated crawlers, focusing 517 on user interactions and website categories. Finally, we compare 518 the prevalence of Javascript APIs in fingerprinting scripts collected 519 520 from real user browsing sessions with that of the automated crawl 521 and discover potential new fingerprinting vectors.

	Scripts (FP / Total)	Websites (FP / Total)
Real Users	695 / 80,969	471 / 3,000
Automated Crawl	210 / 32,823	363 / 3,000

Table 1: Number of scripts and websites detected as fingerprinting by real user browsing session vs automated crawl

4.1 Prevalence of Fingerprinting

To compare against an automated crawl, we follow the same strategy as Annamalai et al. [8], using an instrumented Chrome browser along with Puppeteer to visit the same 3K top-ranked websites provided to the participants of our study. We also try to simulate some degree of user interaction (i.e., scrolling and taking a full-page screenshot) in the automated crawl, but the crawler does not visit any subpage linked from the main website.

In Table 1, we report the number of fingerprinting scripts collected during the study with real users vs. the automated crawl, along with the number of websites that load these fingerprinting scripts. We also defer a more comprehensive breakdown of the fingerprinting scripts identified to Appendix C.

First, we detect 2.3x more fingerprinting scripts through realuser browsing sessions than through the automated crawl. This suggests the latter may indeed miss many fingerprinting scripts in practice, corroborating preliminary findings by Annamalai et al. [8], who show similar results from a small set of websites. However, this increase may be due to the overall increase in the number of scripts collected from real-user browsing sessions on the same set of websites; specifically, we observe a 146% increase in the number of scripts (80.9k vs. 32.8k) compared to a 34% increase in the number of scripts that are fingerprinting (i.e., 0.86% vs. 0.64%). Nevertheless, we observe a significant increase in the percentage of fingerprinting websites; specifically, out of the same 3K websites, we identify 108 more fingerprinting websites (a 30% increase) from real-user browsing sessions than from automated crawlers.

We also find that real-user browsing sessions missed 117 fingerprinting websites that were instead detected from automated crawls. However, this was due to a bug in the implementation of our Chrome extension that prevented the instrumentation script from being injected into iframes, which are often used for browser fingerprinting. In other words, this was an isolated bug and not an inherent limitation of distributing the deployment of browser fingerprinting detection to end users.

Taking the bug into account, the automated crawl missed 225 fingerprinting websites that were identified from real user browsing sessions. This confirms that automated crawlers cannot replicate real user interactions, which can result in the prevalence of fingerprinting being potentially underestimated in practice.

4.2 Undetected Fingerprinting on Websites (by Sub-Pages)

Next, we take a closer look at the user interactions that specifically trigger fingerprinting scripts on websites. To that end, we group the 225 fingerprinting websites undetected by the automated crawler

566

567

580

546

526

582

583

584

585

586

587

588

589

590

591

592

593

594

595

596

597

598

599

600

601

602

603

604

605

606

607

608

609

610

630

631

632

633

634

635

638

Total	Home	Content	Auth	Failed
225	117	63	15	30

Table 2: Number of fingerprinting websites undetected with automated crawl broken down by reason.

by the specific sub-pages that loaded fingerprinting scripts and report this in Table 2.

Failed Visits. First, we find that the automated crawler failed to visit 30 out of the 3K websites (1%). We believe this is due to bot detectors running before the page is loaded; in fact, the majority of these websites (66.7%) return 4XX errors that can be associated with bot detection scripts [8].

Authentication & Content Pages. We find that a non-negligible number of fingerprinting scripts appear on authentication pages (e.g., login, account, and sign-up pages), and the automated crawler only visited 17 such pages as opposed to the 240 visited by real users. As a result, the automated crawler missed 15 websites that were deploying fingerprinting scripts on the authentication page.

Similarly, as the automated crawler did not visit any inner content pages, it missed another 63 websites that were fingerprinting on the content page. While recent work [51] shows that it is possible for automated crawlers to infer whether a visited page is an authentication page using ML-based methods, we note that such pages might themselves deploy bot detection scripts preventing automated crawlers from successfully visiting them. In fact, our automated crawler failed to visit 6 such authentication pages (26.1%).

Home Pages. The automated crawl also missed fingerprinting 611 612 scripts on 117 home pages. This is surprising as these scripts were 613 not triggered even though the automated crawler successfully visited these pages. This indicates that in some cases, even if the web-614 615 site is loaded successfully, specific user interactions are required to 616 trigger fingerprinting scripts. To investigate this more deeply, we 617 manually visited 11 of these websites (a 10% sample) to observe the 618 specific user interactions that triggered the fingerprinting scripts. 619 We found that 5 out of the 11 websites deployed bot detection scripts after the websites were loaded (as opposed to before, which we pre-620 viously categorized as a "failed visit"). Popular vendors for this kind 621 622 of script include PerimeterX¹ (now called Human) and SiftScience². 623 Therefore, the fingerprinting scripts were no longer triggered be-624 cause the automated crawlers were detected as bots. Also note that 625 four websites only start fingerprinting after user consent is received 626 from a cookie consent banner, consistent with findings from prior 627 work [46]. For the remaining two pages, we were unable to pinpoint 628 the exact user interaction that triggered browser fingerprinting, as 629 we could not consistently get the fingerprinting script to load.

In summary, even if automated crawlers could crawl the Web more deeply and mimic certain user interactions, they are often caught out by advanced bot detectors and may not be able to trigger fingerprinting scripts that require specific user interactions.



²https://sift.com/

Figure 2: Percentage of fingerprinting websites for each website category.

Undetected Fingerprinting on Websites (by 4.3 Category)

Next, we analyze the prevalence of fingerprinting websites by category³ in real-user browsing sessions, as presented in Figure 2. We also quantify the prevalence of fingerprinting by website category in the automated crawl; however, due to space limitations, we omit the results here, as we observe no significant difference between the general prevalence of fingerprinting in real-user browsing sessions and the automated crawl. As opposed to prior work [27, 31], we find that E-Commerce, Shopping, and Society (& Lifestyle) categories have much higher rates of fingerprinting than News websites. We believe one of the main reasons for this discrepancy is that previous studies used the discontinued Amazon Alexa ranking [7], whereas we use the CrUX ranking, which is maintained and known to reflect real-user browsing patterns more accurately [50]. Additionally, as no significant difference was observed in the prevalence of fingerprinting by website category between our automated crawl and the real-user browsing sessions, we do not believe this discrepancy was due to real user behaviors. Note that other categories (e.g., adult content) are also not represented in our results, as we had filtered out these websites for privacy reasons.

Finally, we observe that fingerprinting is more likely to be undetected by the automated crawler for a few specific categories of websites than others. To measure this, we introduce the notion of Miss Percentage, defined as the percentage of websites detected as fingerprinting by the real-user browsing sessions but undetected by the automated crawl. Figure 3 shows that Technology and Video Streaming fingerprinting websites are the most likely to be missed by automated crawlers than by real-user browsing sessions. This is expected as Technology and Video Streaming websites might require user interaction or user login before entering the "main page" of the application, where fingerprinting is expected to happen.

4.4 Comparison of Fingerprinting APIs

Next, we investigate the differences in the Javascript APIs used frequently by the fingerprinting scripts captured from real-user browsing sessions, as compared to those captured from automated crawls. To do so, we quantify the relative prevalence of APIs used for fingerprinting following prior work by Iqbal et al. [31] and compute the Call Ratio for each API found in the real-user browsing sessions, i.e., the ratio between the number of times a given API is called by



Anon

websites 40 of FP v 20 Percentage 10 Streamin Entertain E-con Conten

³We categorize the websites using the Cloudflare Domain Intelligence API [15].



Figure 3: Percentage of fingerprinting websites undetected by the automated crawler in each website category.

Javascript API	Call Ratio
audiocontext.sinkid	00
audiocontext.onsinkchange	~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
rtcpeerconnection.getconfiguration	~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
rtcpeerconnection.sctp	~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
svgtextcontentelement.getextentofchar	~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
rtcpeerconnection.tostring	∞
window.navigator.plugins[chrome pdf plugin]	7.50
window.navigator.plugins[webkit built-in pdf]	6.77
window.navigator.plugins[microsoft edge pdf viewer]	6.77
window.navigator.plugins[chrome pdf viewer]	6.77
window.navigator.plugins[chromium pdf viewer]	6.77
window.navigator.plugins[pdf viewer]	5.65
offlineaudiocontext.hasownproperty	3.62

Table 3: Call Ratio of a sample of Javascript APIs predominantly used by fingerprinting scripts in real-world browsing sessions.

fingerprinting scripts and by non-fingerprinting scripts. The Call Ratio is high (above 1) when a given API keyword is used more prevalently in fingerprinting than non-fingerprinting scripts.

Specifically, we use it to identify potential fingerprinting vectors from real-user browsing sessions that automated crawlers might otherwise miss. In Table 3, we report the APIs with high Call Ratios in the real-user browsing sessions that were simultaneously not used by any fingerprinting scripts captured in the automated crawl (i.e., the API would have been missed by the automated crawl). The call Ratio is ∞ when no non-fingerprinting scripts use the keyword.

We observe that audio and WebRTC fingerprinting are more 740 741 prevalent in real-user browsing sessions than automated crawls. 742 Specifically, we observe that multiple Audio APIs (i.e., audiocon-743 text.sinkid, audiocontext.onsinkchange) and WebRTC APIs (i.e., 744 rtcpeerconnection.getconfiguration, rtcpeerconnection.sctp, rtcpeer-745 connection.tostring) are exclusively used by fingerprinting scripts in real-user browsing sessions (Call Ratio = ∞). At the same time, 746 these APIs did not appear to be used by any fingerprinting script in 747 748 the automated crawl. This is probably due to these fingerprinting techniques only occurring if audio devices or peers are present in 749 the network and triggered, which requires real-user devices with 750 an audio interface or a crawler setup that simulates them effectively. 751 752 Nevertheless, prior work [27] has identified the AudioContext and 753 RTCPeerConnection APIs as prominent fingerprinting vectors.

On the other hand, accesses to the Navigator API has previously not been considered a robust signal to detect fingerprinting [27, 31] as they are often used by non-fingerprinting scripts as well. Conversely, our analysis of real-user browsing sessions shows that specific attributes of the Navigator API can, in fact, be used as reliable signals. Specifically, from Table 3, we note that the PDF viewer plugin is used predominantly by fingerprinting scripts to identify the specific browser being used. Upon closer inspection, we observe that fingerprinting scripts use the "length" and "description" attributes of these plugins to identify which PDF viewer is being used by the browser. Although this API is now deprecated, it is still available in all modern browsers [21], and we believe that the differences in how the specification is implemented across different versions of browsers (e.g., returning a hard-coded list) might be a useful fingerprinting vector.

Overall, this confirms that real-user browsing sessions shed better light on fingerprinting vectors used in the wild compared to automated crawlers.

5 Privacy-Preserving Federated Browser Fingerprinting Detection

In the previous section, we showed that automated crawls might indeed miss fingerprinting scripts that are instead captured by real-user browsing sessions. This sheds light on the benefits of training fingerprinting detection models on telemetry from realuser browsing sessions. However, having users share this data with a third party would be both inefficient and extremely privacyinvasive. On the other hand, only training locally (i.e., each user trains their own local model) would yield low accuracy and result in losing crucial knowledge about new fingerprinting behavior.

Recently, Annamalai et al. [8] introduced FP-Fed, a system allowing participants to collaboratively build a browser fingerprinting detection model while keeping their data private, but only sharing model updates using the Federated Learning paradigm [37]. However, they ultimately only tested it on an automated crawl (i.e., they did not capture real user interactions). In the rest of this section, we investigate whether FP-Fed is effective at training browser fingerprinting detectors on real-user datasets.

5.1 Overview of FP-Fed [8]

In a nutshell, FP-Fed works as follows. First, participants train local models to detect browser fingerprinting based on the data collected from their individual browsing sessions. The participants then share the model updates (but not the raw collected data) with a central server, which aggregates the model updates and adds noise to satisfy differential privacy and protect participants' privacy. This differentially private "global" model is then shared with the participants, and the process repeats over multiple rounds until the global model converges. Please see Appendix A for more details.

Improvements. In our work, we do not only deploy FP-Fed but also modify it to improve model performance. Specifically, instead of training the local models from scratch, as done in [8], we first pre-train the local models, non-privately, on public data and then fine-tune them on the private local training datasets. This is a popular approach used to privately train ML models that yields significantly better model performance [19, 53].

734

735

736

737

738

739

754

697

698

699

700

701

702

703

704

705

811

812

755

756

757

758

759

760

761

762

763

764

765

766

767

768

769

770

771

772

773

774

775

776

777

778

779

780

781

782

FP-Fed	on Real-User S	essions	Automated
$\varepsilon = 1$	$\varepsilon = 5$	$\varepsilon = 10$	(Centralized)
0.979 ± 0.013	0.982 ± 0.013	0.983 ± 0.011	0.971 ± 0.021

Table 4: AUPRC of Fed-FP [8] at various levels of privacy (ε) on real-user browsing sessions vs. a centralized model built on the automated crawl.

In our experimental evaluation, we first pre-train the local models on an equivalent automated crawl of the 3K top-ranked websites. As this is public data, pre-training local models on this data can be done non-privately without degrading the privacy guarantees provided by differential privacy. Furthermore, by combining the automated crawl and real-user browser sessions, we argue that the model can achieve precision that is close to models built only on the automated crawl from prior work [31], while learning from many more fingerprinting scripts present in the real-user browsing sessions thus improving the model's recall. We test this claim empirically next.

5.2 Evaluation

813

814

815

816

817

818

819

820

821

822

823

824

825

826

827

828

829

830

831

832

833

834

835

870

We now analyze the feasibility and effectiveness of training de-836 tection models for browser fingerprinting using FP-Fed [8]. To do 837 so, we first partition the scripts from real-user browsing sessions 838 into 80% for training (64,776 scripts, 556 fingerprinting) and 20% 839 for testing (16,193 scripts, 139 fingerprinting). Next, we re-sample 840 the training data using the fine-grained Tranco [47] ranking as 841 done in [8] to simulate 1 million users participating in FP-Fed, as 842 expected in a real-world scenario. Finally, we train a Logistic Re-843 844 gression classifier using FP-Fed at various privacy levels and test 845 the model performance on the test set. As mentioned above, recall that we first pre-train the local models of participants with the 846 847 scripts collected from the automated crawl.

848 In Table 4, we report the Area Under Precision-Recall Curve (AUPRC) statistic along with the standard deviation over 5-fold 849 stratified cross-validation to assess model performance. We choose 850 AUPRC as it summarizes the model performance over many possi-851 ble thresholds that can be tuned to achieve a desired precision/recall 852 trade-off. Furthermore, to assess the robustness of the model in a 853 potential real-world deployment scenario, we compare its perfor-854 855 mance in two distinct cases. In the first one, we start by training a base model (centrally, no federation) with data from the auto-856 857 mated crawl and then fine-tune it with FP-Fed on actual data from 858 real-user browsing sessions. In the second case, which corresponds 859 to the traditional approach commonly used in the literature, we train the central model with data from the automated crawl alone. 860 861 To provide a realistic performance assessment on real-world data, we evaluate both models on the 20% testing set of the real-user 862 863 browsing sessions.

Our results show that even in high privacy settings ($\varepsilon = 1.0$), the model trained on real-user browsing sessions reaches an AUPRC = 0.979, with performance slightly exceeding the centralized model trained with no privacy on the automated crawl only, which achieves AUPRC = 0.971. This is notable considering that the noise introduced by DP makes it challenging for private classifiers to achieve even comparable performance to their non-private counterparts [19]. In fact, even at a moderate privacy level ($\varepsilon = 5.0$), the model trained on real-user browsing sessions outperforms the automated crawl by more than one percentage point.

This confirms that the federated model can leverage the fingerprinting behaviors of the scripts present in the combined browsing sessions from real users that are otherwise unavailable to models trained on the automated crawls alone. While this result is promising, we note that the standard deviations we observe are relatively large, which calls for further research covering a larger number of websites and users.

6 Conclusion

This paper presented the design and execution of a user study geared to investigate the differences in the prevalence and distribution of browser fingerprinting in real-user browsing sessions as opposed to automated crawlers predominantly used by prior work in browser fingerprinting [27, 31]. To do so, we built a Chrome extension and collected fingerprinting scripts from 30 participants as they browsed 3,000 top-ranked websites. We compared the resulting differences by simultaneously performing an automated crawl of the same websites. Additionally, we evaluated the feasibility and effectiveness of collaboratively and privately training a distributed browser fingerprinting detection model using federated learning.

Our analysis showed not only that automated crawls missed a non-negligible amount of fingerprinting scripts but that they also heavily underestimated the prevalence of browser fingerprinting in top-ranked websites. Specifically, we observed that 47.8% of websites identified as fingerprinting from real-user browsing sessions were undetected by the automated crawl. Our findings empirically validate existing hypotheses that discrepancies arise because automated crawlers lack the behavioral nuances of human users. Consequently, bot detection scripts often block them or fail to trigger website fingerprinting mechanisms. Finally, we showed that ML models trained on a combination of crawled data and subsequently fine-tuned in a privacy-preserving way with data from real browsing sessions on-device likely yield better performance on real-world datasets than models trained purely on crawled data. Limitations. Naturally, our work is not without limitations. For instance, the scope of website coverage in our study is limited. While our analysis encompasses an order of magnitude more websites than previous research [8], it represents a subset of the websites visited daily by users. This limited scope was deliberate and focused on protecting the privacy of our participants, as ultimately, the study's primary objective was to investigate discrepancies in fingerprinting prevalence and distribution between real-user browsing sessions and automated crawls.

Having established a demonstrable difference, future research should prioritize expanding the study's reach by increasing both the number of websites analyzed and the participant pool. Crucially, this expansion will necessitate further development of privacypreserving data collection and analysis techniques, a direction we intend to explore in future work.

References

Gunes Acar, Christian Eubank, Steven Englehardt, Marc Juarez, Arvind Narayanan, and Claudia Diaz. 2014. The Web Never Forgets: Persistent Tracking

TheWebConf'25, April 28-May 2, 2025, Sydney, Australia

987

988

989

990

991

992

993

994

995

996

997

998

999

1000

1001

1002

1003

1004

1005

1006

1007

1008

1009

1010

1011

1012

1013

1014

1015

1016

1017

1018

1019

1020

1021

1022

1023

1024

1025

1026

1027

1028

1029

1030

1031

1032

1033

1034

1035

1036

1037

1038

1039

1040

1041

1042

1043

1044

Mechanisms in the Wild. In ACM CCS.

929

930

931

932

933

934

935

936

937

938

939

940

941

942

943

944

945

946

947

948

949

950

951

952

953

954

955

956

957

958

959

960

961

962

963

964

965

966

967

968

969

970

971

972

973

974

975

976

977

978

979

980

981

982

983

984

985

986

- [2] Gunes Acar, Marc Juarez, Nick Nikiforakis, Claudia Diaz, Seda Gürses, Frank Piessens, and Bart Preneel. 2013. FPDetective: Dusting the Web for Fingerprinters. In ACM CCS.
- [3] Mohammed Adnan, Shivam Kalra, Jesse C Cresswell, Graham W Taylor, and Hamid R Tizhoosh. 2022. Federated learning and differential privacy for medical image analysis. *Scientific reports* 12, 1 (2022).
- [4] Seyed Ali Akhavani, Jordan Jueckstock, Junhua Su, Alexandros Kapravelos, Engin Kirda, and Long Lu. 2021. Browserprint: An Analysis of the Impact of Browser Features on Fingerprintability and Web Privacy. In *Information Security* Conference.
- [5] Furkan Alaca and Paul C Van Oorschot. 2016. Device fingerprinting for augmenting web authentication: classification and analysis of methods. In ACM CCS.
- [6] Amazon. 2005. Amazon Mechanical Turk. https://www.mturk.com/.
- [7] Amazon. 2022. We will be retiring Alexa.com on May 1, 2022. https://web.archive.org/web/20220102200605/https://support.alexa.com/hc/enus/articles/4410503838999.
- [8] Meenatchi Sundaram Muthu Selva Annamalai, Igor Bilogrevic, and Emiliano De Cristofaro. 2024. FP-Fed: Privacy-Preserving Federated Detection of Browser Fingerprinting. NDSS.
- [9] Apple. 2023. Apple announces powerful new privacy and security features. https://www.apple.com/sg/newsroom/2023/06/apple-announcespowerful-new-privacy-and-security-features/.
- [10] Pouneh Nikkhah Bahrami, Umar Iqbal, and Zubair Shafiq. 2022. FP-Radar: Longitudinal Measurement and Early Detection of Browser Fingerprinting. PETS (2022).
- [11] Hristo Bojinov, Yan Michalevsky, Gabi Nakibly, and Dan Boneh. 2014. Mobile Device Identification via Sensor Fingerprinting. arXiv:1408.1416 (2014).
- [12] Brave. 2024. Fingerprinting Protections. https://github.com/brave/bravebrowser/wiki/Fingerprinting-Protections.
- [13] Yinzhi Cao, Song Li, and Erik Wijmans. 2017. (Cross-)Browser Fingerprinting via OS and Hardware Level Features. In NDSS.
- [14] Chrome. 2024. Chrome User Experience Report. https://developer.chrome.com/ docs/crux/.
- [15] Cloudflare. 2024. Get Domain Details. https://developers.cloudflare.com/api/ operations/domain-intelligence-get-domain-details.
- [16] Low Pay Commission. 2024. National Minimum Wage in 2024. https://www.gov. uk/government/publications/the-national-minimum-wage-in-2024.
- [17] Luke Crouch. 2018. Improving privacy without breaking the web. https://blog. mozilla.org/data/2018/01/26/improving-privacy-without-breaking-the-web/.
- [18] Anupam Das, Gunes Acar, Nikita Borisov, and Amogh Pradeep. 2018. The Web's Sixth Sense: A Study of Scripts Accessing Smartphone Sensors. In ACM CCS.
- [19] Soham De, Leonard Berrada, Jamie Hayes, Samuel L Smith, and Borja Balle. 2022. Unlocking High-Accuracy Differentially Private Image Classification through Scale. arXiv:2204.13650 (2022).
- [20] Disconnect. 2018. Disconnect defends the digital you. https://disconnect.me.[21] MDN Web Docs. 2024. Navigator: plugins property. https://developer.mozilla.
- org/en-US/docs/Web/API/Navigator/plugins.
 [22] Cynthia Dwork and Aaron Roth. 2014. The Algorithmic Foundations of Differential Privacy. Foundations and Trends in Theoretical Computer Science (2014).
- [23] EasyList. 2024. EasyPrivacy. https://easylist.to/easylist/easyprivacy.txt.
- [24] Peter Eckersley. 2010. How Unique Is Your Web Browser?. In PETS.
- [25] EFF. 2023. PrivacyBadger. https://github.com/EFForg/privacybadgerfirefox/blob/ master/data/cookieblocklist.txt.
- [26] Steven Englehardt. 2020. Firefox 72 blocks third-party fingerprinting resources. https://blog.mozilla.org/security/2020/01/07/firefox-72-fingerprinting/.
- [27] Steven Englehardt and Arvind Narayanan. 2016. Online tracking: A 1-million-site Measurement and Analysis. In ACM CCS.
- [28] U Huws, S Dench, and Ron Iphofen. 2004. An EU code of ethics for socio-economic research. Institute for Employment Studies.
- [29] Muhammad Ikram, Hassan Jameel Asghar, Mohamed Ali Kaafar, Balachander Krishnamurthy, and Anirban Mahanti. 2017. Towards Seamless Tracking-Free Web: Improved Detection of Trackers via One-class Learning. *PETS* (2017).
- [30] Iovation. 2019. Iovation Fraud Protection. https://web.archive.org/web/ 20191130164107/https://www.iovation.com/fraudforce-fraud-detectionprevention.
- [31] Umar Iqbal, Steven Englehardt, and Zubair Shafiq. 2021. Fingerprinting the fingerprinters: Learning to detect browser fingerprinting behaviors. In IEEE S&P.
- [32] Peter Kairouz, Ziyu Liu, and Thomas Steinke. 2021. The Distributed Discrete Gaussian Mechanism for Federated Learning with Secure Aggregation. In *ICML*.
- [33] Pierre Laperdrix, Gildas Avoine, Benoit Baudry, and Nick Nikiforakis. 2019. Morellian analysis for browsers: Making web authentication stronger with canvas fingerprinting. In Detection of Intrusions and Malware, and Vulnerability Assessment.
- [34] Zengrui Liu, Jimmy Dani, Shujiang Wu, Yinzhi Cao, and Nitesh Saxena. 2024. Identified-and-Targeted: The First Early Evidence of the Privacy-Invasive Use of

Browser Fingerprinting for Online Tracking. *arXiv:2409.15656* (2024).

- [35] Ingrid Lunden. 2018. Relx acquires ThreatMetrix for \$817M to ramp up in riskbased authentication. https://techcrunch.com/2018/01/29/relx-threatmetrix-riskauthentication-lexisnexis/?guccounter=1.
- [36] Jonathan R Mayer. 2009. "Any person... a pamphleteer": Internet Anonymity in the Age of Web 2.0. Undergraduate Senior Thesis, Princeton University 85 (2009).
- [37] H Brendan McMahan, Eider Moore, Daniel Ramage, and Blaise Agüera y Arcas. 2016. Federated learning of deep networks using model averaging. arXiv:1602.05629 (2016).
- [38] H Brendan McMahan, Daniel Ramage, Kunal Talwar, and Li Zhang. 2018. Learning Differentially Private Recurrent Language Models. In *ICLR*.
- [39] Luca Melis, Congzheng Song, Emiliano De Cristofaro, and Vitaly Shmatikov. 2019. Exploiting Unintended Feature Leakage in Collaborative Learning. In IEEE S&P.
- [40] Keaton Mowery and Hovav Shacham. 2012. Pixel perfect: Fingerprinting canvas in HTML5. W2SP (2012).
- [41] Mozilla. 2023. WebRTC API. https://developer.mozilla.org/en-US/docs/Web/API/ WebRTC_API.
- [42] Mohammad Naseri, Yufei Han, Enrico Mariconti, Yun Shen, Gianluca Stringhini, and Emiliano De Cristofaro. 2022. Cerberus: Exploring Federated Prediction of Security Events. In ACM CCS.
- [43] Ray Ngan, Surya Konkimalla, and Zubair Shafiq. 2022. Nowhere to Hide: Detecting Obfuscated Fingerprinting Scripts. arXiv:2206.13599 (2022).
- [44] Department of Labor. 2024. State Minimum Wage Laws. https://www.dol.gov/ agencies/whd/minimum-wage/state.
- [45] Łukasz Olejnik, Gunes Acar, Claude Castelluccia, and Claudia Diaz. 2016. The leaking battery: A privacy analysis of the HTML5 Battery Status API. In DPM.
- [46] Emmanouil Papadogiannakis, Panagiotis Papadopoulos, Nicolas Kourtellis, and Evangelos P Markatos. 2021. User Tracking in the Post-cookie Era: How Websites Bypass GDPR Consent to Track Users. In WWW.
- [47] Victor Le Pochat, Tom van Goethem, Samaneh Tajalizadehkhoob, Maciej Korczynski, and Wouter Joosen. 2019. Tranco: A Research-Oriented Top Sites Ranking Hardened Against Manipulation. In NDSS.
- [48] Gaston Pugliese, Christian Riess, Freya Gassmann, and Zinaida Benenson. 2020. Long-Term Observation on Browser Fingerprinting: Users' Trackability and Perspective. *PETS* (2020).
- [49] Franziska Roesner, Tadayoshi Kohno, and David Wetherall. 2012. Detecting and Defending Against Third-Party Tracking on the Web. In USENIX.
- [50] Kimberly Ruth, Deepak Kumar, Brandon Wang, Luke Valenta, and Zakir Durumeric. 2022. Toppling Top Lists: Evaluating the Accuracy of Popular Website Lists. In ACM IMC.
- [51] Asuman Senol, Alisha Ukani, Dylan Cutler, and Igor Bilogrevic. 2024. The Double Edged Sword: Identifying Authentication Pages and their Fingerprinting Behavior. In ACM WWW.
- [52] Lichao Sun, Jianwei Qian, and Xun Chen. 2021. LDP-FL: Practical Private Aggregation in Federated Learning with Local Differential Privacy. In *IJCAI*.
- [53] Florian Tramer and Dan Boneh. 2021. Differentially Private Learning Needs Better Features (or Much More Data). In *ICLR*.
- [54] Stacey Truex, Ling Liu, Ka-Ho Chow, Mehmet Emre Gursoy, and Wenqi Wei. 2020. LDP-Fed: Federated learning with local differential privacy. In ACM International Workshop on Edge Systems, Analytics and Networking.
- [55] W3C. 2021. Mitigating Browser Fingerprinting in Web Specifications. https: //w3c.github.io/fingerprinting-guidance/.

A FP-Fed [8]

FP-Fed [8] is a privacy-preserving distributed system for detecting browser fingerprinting in the wild. Using Differentially Private Federated Learning (DP-FL) [37], FP-Fed enables learning the fingerprinting behaviors of scripts from the combined browsing sessions of many users while providing strong formal guarantees of privacy. There are five main steps to FP-Fed, also depicted in Figure 4:

- (1) Participants run an instrumented browser (e.g., with a Chrome extension installed) and naturally browse websites.
- (2) The instrumented browser / extension collects script execution data and performs the following actions before federated training begins:
 - (a) Collects number of times specific (often high-entropy) monitored APIs are called, along with associated return values and arguments;



Figure 4: An overview of FP-Fed [8]: (1) Participants browse websites and collect script execution data. (2) For each round of training, the server sends the previous round's global model parameters to a few selected participants. (3) These selected participants train a local model based on the collected local script execution data, and (4) send the local model updates to the server, (5) which aggregates them and adds differentially private noise to the aggregate to form the global model of the current round. Steps (2) to (5) are then repeated until the model converges.

- (b) Extracts high-level features for each script loaded on any visited website (e.g., length of string argument, number of values in list returned, etc.);
- (c) Assigns seed ground truth labels (fingerprinting/nonfingerprinting) to each script, according to a highprecision heuristic;
- (d) Participates in a differentially private federated feature pre-processing phase that normalizes extracted highlevel features in a privacy-preserving way.
- (3) At each round, the server sends the previous round's global model parameters to a subset of participants selected.
- (4) Participants initialize a local model with the previous round's global model parameters and train their local model with their locally collected data.
- (5) Participants send the model updates to the server.
- (6) The server aggregates the model updates, adds statistical noise, and computes the global model parameters for the next round.
- (7) Steps 2) to 5) are repeated over multiple rounds until model convergence;

Once the global model converges in FP-Fed, training stops, and all participants can use the trained global model for on-device browser fingerprinting detection by all participants. FP-Fed can be run regularly (e.g., once every few weeks) so that the model can be updated to learn the latest fingerprinting behaviors.

Type of FP	Real User	Automated
Canvas	629	177
Canvas Font	40	24
Audio	215	107
WebRTC	85	35
Total	695	210

Table 5: Number of fingerprinting scripts captured by real user browsing session vs. automated crawls. *NB*: The same script can perform multiple types of fingerprinting.

B Human Intelligence Task (HIT)

In Figure 5, we display the Human Intelligence Task (HIT), i.e., the ad, used to recruit participants from the Amazon mTurk platform. Specifically, prospective participants were provided with some task instructions, an information sheet describing in layman terms the purpose of the data collection and the types of data being collected and a consent form to ensure they have read and understood their data rights and task requirements. They were given a password to authenticate themselves with the Chrome extension used for data collection and were told to enter the "Task Completion Code", which was provided to them by the extension once they completed browsing the assigned websites. For brevity, a significant portion of the consent form has been left out from Figure 5.

Anon

1161 C Prevalence of Types of Fingerprinting

In Table 5, we report the number of each type of fingerprinting
 script captured by real user browsing sessions and compare it with
 the scripts collected from the automated crawl.

Overall, similar to the general prevalence of fingerprinting scripts, we note no major difference between the presence of each type of fingerprinting. Specifically, in both the real user browsing sessions and the automated crawl, Canvas FP is observed to be the most popular form of fingerprinting, followed by Audio, WebRTC, and Canvas Font.

1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1250
1230
1250
1250 1251 1252
1250 1251 1252 1253
1250 1251 1252 1253 1254
1250 1251 1252 1253 1254 1255
1250 1251 1252 1253 1254 1255 1256
1250 1251 1252 1253 1254 1255 1256 1257
1250 1251 1252 1253 1254 1255 1256 1256 1257 1258
1250 1251 1252 1253 1254 1255 1256 1257 1258 1259
1250 1251 1252 1253 1254 1255 1256 1257 1258 1259 1260
1250 1251 1252 1253 1254 1255 1256 1257 1258 1259 1260 1261
1250 1251 1252 1253 1254 1255 1256 1257 1258 1259 1260 1261 1262
1250 1251 1252 1253 1254 1255 1255 1256 1257 1258 1259 1260 1261 1262 1263
1250 1251 1252 1253 1254 1255 1256 1257 1258 1259 1260 1261 1262 1263 1264
1250 1251 1252 1253 1254 1255 1256 1257 1258 1259 1260 1261 1262 1263 1264 1264
1250 1251 1252 1253 1254 1255 1256 1257 1258 1259 1260 1261 1262 1263 1264 1265 1266
1250 1251 1252 1253 1254 1255 1256 1257 1258 1259 1260 1261 1262 1263 1264 1265 1266 1267
1250 1251 1252 1253 1254 1255 1256 1257 1258 1259 1260 1261 1262 1263 1264 1265 1266 1267 1268
1250 1251 1252 1253 1254 1255 1256 1257 1258 1259 1260 1261 1262 1263 1264 1265 1266 1267 1268 1269
1250 1251 1252 1253 1254 1255 1256 1257 1258 1259 1260 1261 1262 1263 1264 1265 1266 1267 1268 1269 1270
1250 1251 1252 1253 1254 1255 1256 1257 1258 1259 1260 1261 1262 1263 1264 1265 1266 1267 1268 1269 1270
1250 1251 1252 1253 1254 1255 1256 1257 1258 1259 1260 1261 1262 1263 1264 1265 1266 1267 1268 1269 1270 1271
1250 1251 1252 1253 1254 1255 1256 1257 1258 1259 1260 1261 1262 1263 1264 1265 1266 1267 1268 1269 1270 1271 1272
1250 1251 1252 1253 1254 1255 1256 1257 1258 1259 1260 1261 1262 1263 1264 1265 1266 1267 1268 1269 1270 1271 1272 1273
1250 1251 1252 1253 1254 1255 1256 1257 1258 1259 1260 1261 1262 1263 1264 1265 1266 1267 1268 1269 1270 1271 1272 1273 1274
1250 1251 1252 1253 1254 1255 1256 1257 1258 1259 1260 1261 1262 1263 1264 1265 1266 1267 1268 1269 1270 1271 1272 1273 1274 1275

TheWebConf'25, April 28-May 2, 2025, Sydney, Australia

1277	Task Instructions (Click to collapse)			
278	We are conducting a study on browser fingerprinting (https://developer.mozilla.org/en-US/docs/Glossary/Fingerprinting), which is an advanced online tracking technique on the web. We want to understand how prevalent this technique is on various popular websites. To that end, we are			Fingerprinting), which is an advanced opular websites. To that end, we are
279	conducting this user study, which you have volunteered to participate in. Please first read the Participation Information Sheet Then fill in the Consent Form below. Next, proceed to install the chrome			Next, proceed to install the chrome
80	extension fre automatical	om the link provided below. Enter the password you find below into the extension and ly downloaded. Next, open that file and you will see a list of websites. Your task is to v	I a file na /isit and	med sites to browse.txt will be interact with each website in that list,
81	By clicking on at least 10 links or buttons you find on each website. Wherever possible, visit login pages (but do not login), accept cookie consents, and solve CAPTCHAs when browsing the websites. After visiting all of the websites, click on the Done browsing all websites button			
2	and you will receive a task completion code to paste into the box below. Lastly, you can proceed to uninstall the extension and you will receive credit for completing the task. If necessary, please find more detailed instructions here			
[Consent Fo	sm		
	Please com	nplete this form after you have read the Information Sheet and/or listened to an e	explanat	ion about the research.
	Title of Stur	dy: Collecting Real World Website Data for Browser Fingerprinting Detection		
	Thank you fe	or considering taking part in this research. The person organising the research must e	explain th	e project to you before you agree to
	take part. If decide whet	f you have any questions arising from the Information Sheet or explanation already give ther to join in. You will be given a copy of this Consent Form to keep and refer to at ar	en to yo ny time.	u, please ask the researcher before you
	I confirm th will be assu	ast I understand that by ticking/initialling each box below I am consenting to this umed that unticked/initialled boxes means that I DO NOT consent to that part of the second second second second	elemen the stud	t of the study. I understand that it ly. I understand that by not giving
	consent for	r any one element that I may be deemed ineligible for the study.		
			-	1
			Tick Box	
		I confirm that I have read and understood the Information Sheet for the		-
	1.	what will be expected of me. I have also had the opportunity to ask questions which have been answered to my satisfaction and would like to take part in the obven data collection - winter		
		Leopent to participate in the study is determined that much	-	-
		I consent to participate in the study. I understand that my personal information, namely the websites I visit (restricted to websites in the list provided to me), and the data from the websites (code executed by		
	2.	websites and browser features accessed by the website restricted to the list provided) will be used for the purposes explained to me. I understand that according to data protection legislation 'performance		
		of a task in the public interest' will be the lawful basis for processing.		_
		Use of the information for this project only		
	3.	that all efforts will be made to ensure I cannot be identified. Lunderstand that my data cathered in this study will be needdonymized		
		and stored securely. It will not be possible to identify me in any publications.		
		I understand that my information may be subject to review by	_	-
	4.	purposes.		
		I understand that my participation is voluntary and that I am free to withdraw at any time within 2 weeks of the task deadline as detailed in		
	5	the MTurk platform without giving a reason. I understand that if I decide to withdraw, any personal data I have		
	0.	provided up to that point will be deleted unless I agree otherwise. After 2 weeks, I understand that it will no longer be possible to remove		
		ary use at a it will have been pseudonymized and the data analysis will already have been conducted.		
	6.	I understand the potential risks of participating and the support that will be available to me should I become distressed during the course of the		
		research.	-	-
	7.	I understand the direct/indirect benefits of participating.		
		I understand that the data will not be made available to any commercial	_	-
	8.	organisations but is solely the responsibility of the researcher(s) undertaking this study.		
	9	I understand that I will be compensated for the portion of time spent in	0	
		the study (if applicable) or fully compensated if I choose to withdraw.	-	-
	10.	I understand that the information I have submitted will be published as a research paper but it will not be possible to identify me in the paper.		
		I harshy confirm that I understand the inclusion mitaging and in the inclusion	-	-
	11.	Interest commitment in uncerstand the inclusion criteria as detailed in the Information Sheet.		
		I hereby confirm that:		
	12.	 a. I understand the exclusion criteria as detailed in the Information Sheet; and b. I do not fall under the exclusion criteria 		
			-	-
	13.	I am aware of who I should contact if I wish to lodge a complaint.		
		I am aware of the use of information for this project and beyond for		1
	14.	I would be happy for the data I provide to be archived at a		
	15	Lucionate il come a tele est in this is in the	-	
	15.	i voluntarniý agrée to take part in this study.		
l				
		Chrome		



Anon.