

# adF: a Novel System for Measuring Web Fingerprinting through Ads

Miguel A Bermejo-Agueda, Patricia Callejo, Rubén Cuevas, Ángel Cuevas, Álvaro Mayol

**Abstract**—This paper introduces *adF*, a novel system for analyzing the vulnerability of different devices, operating systems, and browsers to web fingerprinting. *adF* performs its measurements from code inserted in ads. We use our system in several ad campaigns that delivered 5.40 million ad impressions. The collected data allow us to assess the vulnerability of current desktop and mobile devices to web fingerprinting. We estimate that 66% of desktop and 40% of mobile devices can be uniquely fingerprinted with our web fingerprinting system. However, web fingerprinting resilience varies significantly across browsers and device types, with Chrome on desktops being the most vulnerable configuration.

To counter web fingerprinting, we propose *ShieldF*, a straightforward solution that blocks browsers from reporting attributes with the most significant discrimination power, as identified through our dataset analysis. Our experiments reveal that *ShieldF* outperforms all anti-fingerprinting solutions proposed by major browsers (Chrome, Safari and Firefox), providing a resilience boost of up to 62% for some device configurations. *ShieldF* is accessible as an add-on for any Chromium-based browser. Moreover, it is readily adoptable by browser and mobile app developers. Its widespread use would lead to a marked enhancement in browsers' and mobile apps' defense against web fingerprinting.

**Index Terms**—fingerprinting, advertising, user privacy, adTag, web tracking, browser developers.

## I. INTRODUCTION

**D**IGITAL advertising is arguably the primary source of funding for the current Internet. Some of the most important tech companies, such as Google or Meta (previously Facebook), declare that over 77.5% and 97.5% (2023 data) of their profits come from digital advertising [1], [2], respectively.

The appeal of online advertising for advertisers, compared to traditional forms of publicity (e.g., TV, radio, or newspapers), lies in its ability to deliver personalized ads. Implementing personalized advertising is based on techniques that identify unique users online. These techniques are used for tracking, profiling users, displaying targeted or re-targeted ads, implementing attribution models, etc. The most widely used techniques for this purpose are third-party cookies, advertising IDs, and fingerprinting [3].

M. Bermejo-Agueda is with the Telematic Engineering Department, Universidad Carlos III of Madrid, 28911 Leganés, Spain, email: mibermej@pa.uc3m.es.

P. Callejo is with the Telematic Engineering Department, Universidad Carlos III of Madrid, 28911 Leganés, Spain, email: pcallejo@it.uc3m.es.

R. Cuevas is with the Telematic Engineering Department, Universidad Carlos III of Madrid, 28911 Leganés, Spain, email: rcuevas@it.uc3m.es.

A. Cuevas is with the Telematic Engineering Department, Universidad Carlos III of Madrid, 28911 Leganés, Spain, email: acrumin@it.uc3m.es.

Á. Mayol is with Taptap Digital, 28020 Madrid, Spain, email: alvaro.mayol@taptapdigital.com.

Third-party cookies and advertising IDs, when properly implemented with consent acquisition (not always the case), comply with data protection regulations. Fingerprinting, by definition, is an intrusive practice and represents a severe threat to end-users' privacy. The most common type of fingerprinting involves embedding a script in a website. When a user visits such a website, the script collects attributes from the browser, OS, and device rendering the web-page. The concatenation of these attributes creates a device fingerprint. If the number of attributes is large enough, the fingerprint can make the device unique even among hundreds of thousands or millions of devices. This practice is known as *web fingerprinting*. The privacy risks associated with web fingerprinting have prompted extensive research in this area [4]–[9]. A body of work proposes countermeasure techniques to prevent fingerprinting [7], [10]–[19]. Browser developers (e.g., Apple, Mozilla, or Google) have also proposed anti-fingerprinting measures in later releases of their browsers [20]–[22]. However, a rapid web technology development has made some existing proposals obsolete and requires a continuous effort from browser developers to update their anti-fingerprinting solutions.

In this rapidly evolving context, proper auditing methodologies are crucial for measuring the risk of web fingerprinting to users and assessing vulnerabilities in different web browsers. The research community has proposed measurement solutions in this direction [5], [7], [9], [23]. Essentially, these solutions involve deploying a script on a web-page to compute the device fingerprint when a device connects to the page. This methodology assesses whether the computed fingerprint is unique within the database of previously computed fingerprints. This approach has been very useful in improving our understanding of web fingerprinting. However, it presents some limitations. On the one hand, if the auditing script is deployed on a website meant for measuring fingerprinting [5], [7], it requires the voluntary participation of users who are willing to assess their fingerprint. Sustaining a continuous flow of visitors to these types of website over time poses a challenge. On the other hand, running the script on a regular website [8], [24], [25] introduces biases in the type of devices/users visiting that single website. Moreover, the ability to conduct such experiments depends on the willingness of the website owner to participate. Furthermore, the coverage of this technique is limited to web-browsers, leaving aside other very common venues such as mobile apps that operate on top of WebViews, essentially utilizing web technologies.

In this paper, we propose a novel measurement technology to study web fingerprinting. Our methodology utilizes ads, which are HTML elements included on web-pages and within

mobile apps. We embed a script to compute the device fingerprint within these ads, overcoming limitations of existing solutions: 1) We run ad campaigns that perform measurements transparently to users, ensuring longitudinal analyses without relying on volunteers; 2) Our methodology allows for the study of web fingerprinting within the context of mobile apps; 3) Our data sample is not limited to measurements from a single vantage point (i.e., web-page); 4) Researchers have full control over our technique, deciding when to run the campaigns; 5) By harnessing the targeting capabilities of online advertising, our methodology enables the definition of specific device configurations (based on device type, OS, and browser) to be measured for the first time.

We implement our methodology in a system we refer to as *adF* (as discussed in Section IV) and conducted real ad campaigns that collected more than 4,04M fingerprint samples in three independent periods of time over the course of a year and a half (as detailed in Section V). First, these real experiments demonstrate that our methodology achieves an accuracy exceeding 80%, confirming its validity for the study of web fingerprinting. Second, the analysis of the collected fingerprint samples allows to assess the vulnerability of different device configurations to web fingerprinting. We define a device configuration as the combination of device type (desktop vs. mobile), an operating system (e.g., Windows, macOS, Android, or iOS), and a browser (e.g., Mozilla Firefox, Chrome, or Safari). In the case of mobile apps, we only distinguish between operating systems, Android vs. iOS (see Section VII).

Our analysis reveals that 66% of desktop devices can be uniquely fingerprinted from ads rendered by browsers. In the case of mobile devices, 40% (35%) can be fingerprinted with ads delivered to browsers (mobile apps). Moreover, we observe varying degrees of vulnerability to fingerprinting among different device configurations. For instance, in mobile devices, iOS offers significantly better protection against web fingerprinting compared to Android. In the case of desktop devices, over 78% of devices running Chrome can be fingerprinted, regardless of the operating system. This percentage drops to around 5.5% for {Safari, macOS} and stands at 22% for {Firefox, Linux}. Finally, a thorough analysis of all attributes contributing to web fingerprinting reveals that cardinality (number of distinct values an attribute can take) and entropy (frequency of appearance of the different values) are the two crucial parameters that define the discrimination power of attributes. The device configurations most vulnerable to fingerprinting are those associated with a higher number of attributes with significant discrimination power (as discussed in Section IX). Building on this insight, we propose *ShieldF* a straightforward approach, easily adoptable by browser developers, to counter web fingerprinting. *ShieldF* is a browser extension compatible with all Chromium-based browsers. It prevents the reporting of attributes with high discrimination power to third-parties while preserving the user's browsing experience. Our experiments show that *ShieldF* outperforms most popular countermeasure defenses in the literature, and the anti-fingerprinting solutions of all major browser developers (Section X), enhancing the resilience to web fingerprinting across all considered device configurations by up to 62%.

## II. BACKGROUND

### A. Identifying individual users

Advertising companies use different techniques to identify unique users and target them based on individual preferences such as user demographics, location, and interests. Although they can generally only identify devices rather than users (as in the open advertising ecosystem of websites and mobile apps, there is no registration process), they assume that a single user uses a device. According to [3], the most common techniques used to identify unique users are: third-party cookies, advertising IDs, and web fingerprinting techniques.

1) *Third-party cookies*: A cookie is a small file placed by third-parties on the browser that serves as a unique identifier for a user. Whenever the browser visits a website where the third-party is present, the third-party receives its cookie and knows the user visited the current website. Using this technique, the third-party company can (partially) reconstruct the users' web browsing history in the specific browser where the cookie is placed. In addition, if the third-party can classify the visited websites into categories, it can infer the person's interests. Third-party cookies are also used for other purposes, such as re-targeting, which allow companies to remind users of their products or services through ads, emails, social media, or other methods.

2) *Advertising ID*: The advertising ID (Ad-ID) is a unique user identifier from the UUID standard assigned to an electronic device (smartphone, tablet, PC) and is intended for advertising purposes. It allows advertisers to track users' advertising activity to personalize offers through browsers and mobile apps. The Ad-ID is embedded in the ad requests received by ad-tech stakeholders, which use the Ad-ID to enrich the information known about such Ad-ID, e.g., demographic information or interests associated to the Ad-ID. The Ad-ID is available in all major desktop and mobile operating systems (Windows, iOS, macOS, Android, etc.) and usually receives a different name for each operating system, for instance, AAID (Android Advertising ID) in Android or IDFA (Identifier for Advertisers) in Apple's operating systems.

3) *Web fingerprinting*: Web browsers have developed a set of APIs that allow accessing different parameters of the browser, the operating system, or the device. The original goal of these functions is to optimize the functionality of browsers and the user experience. For instance, one of these functions exposes the device's screen size such that the webpage's content can be appropriate for the specific screen size.

Web fingerprinting is implemented through a script embedded in the HTML code of the web-page. This script leverages the different functions and APIs offered by the browser to collect different attributes from the browser configuration (e.g., browser add-ons or configured languages), the OS (installed fonts), and the device (sound device or graphic card information). In conjunction, the values of all these attributes create a fingerprint of the device. If a sufficiently large number of attributes are collected, the resulting fingerprint of the device might be unique even among a large pool of devices.

Note that, contrary to the case of third-party cookies or advertising IDs, which require implicit or explicit consent from

the user, fingerprinting is implemented without any type of consent from the user and thus it is the most privacy-aggressive practice of the three. Nonetheless, web fingerprinting also has legit uses, such as pursuing fraud or credential theft, by verifying the legitimacy of users when logging into specific sites [26].

### B. Web measurements from ads

The online advertising industry has evolved into an intricate ecosystem for measuring user engagement, drawing on data from JavaScript scripts embedded within ads delivered to websites and mobile apps. These scripts, belonging to different third-party companies, exploit browser APIs to collect data, which is then processed and sent to third-parties' back-end servers. These scripts can be used to measure KPIs (e.g., the click-through ratio of an ad campaign), detect fraudulent practices, and more. In this paper, we use this same measurement technology to implement our script to study web fingerprinting. We embed a script in an ad (rather than a webpage) and gather information about the browser configuration, OS, and device software and hardware while aiming to identify a user uniquely. Nonetheless, note scripts executed within ads face limitations compared to those running natively in a website's HTML code due to security measures like the Cross-Domain Policy.

1) *Web fingerprinting from ads*: One of the main uses of web fingerprinting is to reconstruct the web browsing history of a user. A tracking entities embeds their fingerprinting script in as many websites as possible to gather data about the user's browsing history. It can also be used for re-targeting, i.e., users are shown ads about products they previously viewed on other websites or e-commerce sites. However, web fingerprinting may not be effective in these tasks since it is not feasible to show the same ad to a user on every page they visit. Nonetheless, web fingerprinting from ads can be very useful in building attribution models for advertisers. An attribution model is an extensively used and very valuable tool used by advertisers to understand the performance of their ad campaigns, web-pages, and ad channels. To build this model, advertisers must reconstruct each user's path. This is a vector including all the ads shown to an individual user and the interaction of the user with them. By analyzing the paths of thousands of users, a model can be constructed to infer the most successful elements (ad creatives, ad campaigns, web-pages, mobile apps, etc) that lead to more conversions, such as purchases.

To create users' paths, it is crucial to uniquely identify each user. In this context, web fingerprinting from ads is a plausible alternative to building attribution models. Even with the imminent discontinuation of third-party cookies [27], and the significant restrictions imposed on the use of the advertising ID [28] which limits the number of devices disclosing it<sup>1</sup>, it may become the only practical approach to building attribution models in the near future.

<sup>1</sup>Just 48,52% of the 4.04M samples collected in the dataset used in this paper reported the Ad-ID.

In summary, web fingerprinting from ads also poses a severe threat to end-users privacy, while there is a clear incentive (attribution models) for advertising stakeholders to (potentially) use it. Our main goal is to assess the vulnerability of different device configurations to web fingerprinting. In the context of this paper, a device configuration is defined by the combination of three elements: device type (mobile vs. desktop), operating system (e.g., Windows or iOS), and browser (e.g., Chrome, Safari, or Firefox). In addition, we also consider the specific case of apps on mobile devices (Android and iOS).

### C. Popular browser' approaches to counter fingerprinting

Due to the threat that fingerprinting represents to users' privacy and the social pressure generated, browser developers have proposed different solutions to counter web fingerprinting. These techniques propose modifying, overriding, randomizing, or blocking the value reported for certain attributes. Next, we briefly describe the solutions defined by different browsers to counter web fingerprinting:

- *Brave*: While Brave is subject to web fingerprinting, it is not exposed to web fingerprinting from ads since the advertisements shown by brave are simple text notifications where code cannot be embedded. To mitigate web fingerprinting, Brave blocks several APIs (such as Canvas, Battery Status, WebGL, and Web Bluetooth, among others). These functions are natively integrated into the browser [29].

- *Tor*: It is the browser proposing the most robust solution against fingerprinting. Its approach is to build the same fingerprint for all devices on the web [30]. However, this approach requires blocking some essential functions, which affects the user's browsing experience [31].

- *Safari*: In 2018, Apple released a version of Safari with an anti-fingerprinting setting already built into the browser and thus requiring no user action [20]. In particular, Safari does not support specific APIs such as Web Bluetooth, Device Memory, or User Permission State (e.g., camera, geolocation, or microphone). It also removes detailed information on the user agent or on the type of graphic card used.

- *Firefox*: Mozilla Firefox offers user-configurable privacy options through fingerprint controls. In 2019, Firefox enabled the Enhanced Tracking Protection feature by default as part of its settings [21]. Today, this functionality offers opt-in fingerprint protection preventing the collection of information related to time zone, installed fonts, window size preferences, and operating system versions (among others).

- *Chrome*: It is the browser with the largest market share. Google announced in August 2019 that they would introduce fingerprint-blocking features in Chrome but without providing an expected release date. It was not until early 2020 that Google revealed their plans: reduce the granularity of the information provided by the user agent, which is sent through every HTTP request by default [22]. This functionality is still in the experimental phase.

- *Edge*: It is a web browser developed by Microsoft based on the Chromium engine. Then, it offers similar levels of protection to fingerprinting as Google Chrome.

- *MiuiBrowser*: This Xiaomi-developed browser is a non-exclusive browser for MIUI. While we have been unable to find any officially declared anti-fingerprint technique for this browser, its advanced configuration options allow the change of the user agent by setting the value to desktop mode, iPhone, or iPad.

All these solutions have been designed to counter web fingerprinting by blocking access to certain APIs or functions from any code run by the web browser. This includes code run from ads. Therefore, they are also valid for the case of web fingerprinting from ads (with the exception of Brave, where web fingerprinting from ads does not apply).

Finally, in the case of mobile apps, operating systems such as iOS and Android are implementing significant changes to increase privacy and security by restricting access to data that could be used for fingerprinting [32]. Based on our results, we find that iOS and Android restrict access to information that might be used for fingerprinting from their mobile applications.

### III. THREAT MODEL FOR THE *adF*

In this section, we detail the specific threat model for *adF*, the system we present in Section IV to evaluate the vulnerability of different device configurations to web fingerprinting.

#### A. System model

Our system model consists of a user who uses a (web) browser, such as Chrome, Safari, Firefox, etc., to browse the Internet. The browser runs on any supported device, including desktops, laptops, smartphones, or tablets, and can be run on any relevant operating system such as Windows, macOS, Linux, iOS, or Android. The user can also use mobile applications that operate on Android or iOS WebView. As users access web-pages or mobile applications, they may encounter both first-party and third-party JavaScript code. The browser or app executes the code using its JavaScript engine, such as V8 in Chrome, Edge, or the Android WebView; JavaScriptCore (also known as Nitro) in Safari and iOS WebView; or SpiderMonkey in Firefox. The execution of this code enables interaction with the browser and device through various supported JavaScript APIs. For instance, it can retrieve the user-agent information, screen dimensions, audio and video settings, etc.

#### B. Attacker model

We consider a *remote* attacker who wants to track a web user through *web fingerprinting* techniques embedded in advertisements. In the threat model, the attacker is a malicious ad-tech provider or any other malicious player with access to advertising platforms capable of running ad campaigns. The attacker can add a JavaScript (JS) pixel into ads, which is widely accepted and not considered an illegal technique in ad-tech. This makes it relatively easy for the attacker to perform the attack without the need for third-party collaboration or black-hat hacking techniques to gain access to any third-party venue (e.g., a web-page). This approach significantly lowers

the barriers to conducting such attacks, allowing potential reach to tens of millions of users for fingerprinting purposes.

In contrast, the attacker model defined in previous works in the literature [5], [7], [8], [23], [24], [33] requires the attacker to gain access to one or multiple web-pages to inject a fingerprinting script. This requires either the cooperation of a third party, such as the web-page owner—who may be hesitant due to data protection concerns (e.g., GDPR compliance)—or illegally gaining access to the website by black-hat hacking techniques. To reach tens of millions of users, the accessed websites should be relatively popular, which (in principle) makes it even harder to reach an agreement with their owners due to potential bad reputation implications or hack them (since it is assumed popular web-pages have better security measures in place compared to unpopular ones). Therefore, traditional threat models presented in the literature depict a much higher barrier to performing a fingerprinting attack compared to the alternative using ads discussed in this paper.

### IV. *adF*: A SYSTEM TO MEASURE WEB FINGERPRINTING FROM ADS

*adF* is a novel system designed to assess the vulnerability of different device configurations to web fingerprinting. Figure 1 shows the scheme of the *adF* system, formed by three components: i) The *Data Collector*, a piece of code embedded in online ads that allows collecting a large number of attributes from the browser or the app rendering the ad; ii) The *Fingerprint Constructor*, receives as input the value of the attributes collected by the Data Collector and uses encoding techniques to create a fingerprint; iii) The *Fingerprint Classifier*, decides whether a fingerprint is a unique fingerprint, and thus uniquely identifies a user, or not. The final output of our system is a device fingerprint, categorized as either unique or non-unique.

#### A. Data Collector

To collect attributes from the device, we develop an *adTag*, a JavaScript (JS) code, which can be embedded into any ad format capable of running a script. Our *adTag* is executed upon the rendering process of the ad on a device. The use of *adTags* is a widespread practice in online advertising. Moreover, the research community has also used *adTags* to analyze transparency aspects of online advertising [34] and to conduct network measurements [35].

Our *adTag* captures a comprehensive set of up to 191 diverse attributes from the device, OS, and browser where the ad is displayed. Specifically, it leverages JavaScript APIs within browsers to retrieve information. While some APIs necessitate permissions for access, such as the microphone, camera, or location, the majority of them are openly accessible through our *adTag*. The attributes retrieved by our *adTag* via these APIs are stored on a server within our domain, where we process them and construct the fingerprint using the Fingerprint Constructor. Table IV in Appendix B provides the complete list of attributes considered, along with their respective sources.

An initial analysis of the 191 considered attributes, based on tens of thousands of collected samples, reveals that some of




Figure 1: *adF* visual representation, our web fingerprinting system.

these attributes are not suitable for creating unique fingerprints. Therefore, our Data Collector does not retrieve attributes that meet any of the following criteria: 1) the attribute reports a constant value (e.g., Bars settings: `menubar`); 2) the value of the attribute is unstable over time (e.g., Network Information: `downlink`); 3) the value is unreliable (e.g., Window navigator: `onLine`); 4) most popular browsers do not support the attribute (e.g., WebHID API). After applying this filtering, the final version of the Data Collector retrieves a total of 66 attributes from the initial 191. These attributes are highlighted in bold in Table IV. Note that some of these 66 attributes are joined to form meta-attributes defined as a concatenation of several individual attributes. Table IV also details the mapping of attributes to meta-attributes.

In the case of mobile apps, just a subset of 35 of these 66 attributes are considered and thus used in our analyses. These are identified in Table IV with an asterisk (\*).

Ads are expected to be exposed on a device for a limited time. Usually, a few seconds [36]. Therefore, our *adTag* must collect the 66 (35) attributes and deliver them to the backend server before the user closes the website (or the mobile app) where the ad is displayed. To this end, we use AJAX (Asynchronous JavaScript and XML), which allows the asynchronous execution of code so that the most time-consuming operations do not block the rest of the script's tasks. By doing so, we ensure that for more than 95% of the ads we received a response from the browser (mobile app) including the complete list of attributes. Note that we discard incomplete fingerprint samples to avoid generating fingerprints that may differ due to these uncollected attributes.

### B. Fingerprint Constructor

This module aims to build a device's fingerprint using the 66 (35) attributes collected by the Data Collector from the browser (mobile app). To this end, the Fingerprint Constructor first generates a string concatenating the value of these 66 (35) attributes. On average, this string has a size of 657 (489) Bytes among the samples available in our dataset (as described in Section V). We could use this string as a fingerprint of the device. However, for enhanced computational efficiency, we apply an MD5 hash function to create an alphanumeric fingerprint consisting of 512 bits for each device. The use of a hash function ensures a very low collision probability,

and the resulting fingerprint's reduced size enables rapid computational processing.

### C. Fingerprint Classifier

The fingerprint delivered by the Fingerprint Constructor is actionable, for instance, to track a device on the web, only if it is unique. Consequently, the final component of our system is a binary (fingerprint) classifier. Its output, i.e., the dependent variable, is a binary indicator that determines whether a fingerprint is unique (1) or non-unique (0). It receives as input (i.e., independent variables) the 66 (35) attributes from the Data Collector for browser (mobile app) samples.

The Fingerprint Classifier, like any classifier, requires training. To achieve this, we use the *fingerprint* datasets, as introduced in Section V, comprising 161k and 52k real fingerprints from browsers and mobile apps, respectively. Each fingerprint in these datasets is labelled as either unique or non-unique.

The algorithm employed to build the model for our Fingerprint Classifier is XGBoost (eXtreme Gradient Boosting), which outperformed alternative methods—such as Logistic Regression, Support Vector Machine, Decision Tree, Random Forest, Gradient Boosting Machine, as well as third-party gradient boosting algorithms like LightGBM and CatBoost—in our evaluations. To robustly estimate model performance while maximizing the use of our labelled data (samples with available Advertising ID, see Section V), we adopt a 10-fold stratified cross-validation procedure—an approach to offer an effective balance between bias and variance for moderate-size datasets [37]. Specifically, the dataset is randomly partitioned into 10 equal-sized folds, where in each iteration, the model is trained on 9 folds and tested on the remaining fold. This approach ensures that every fingerprint is used once for testing without any overlap between training and validation, thereby preventing potential data leakage.

Two essential tasks related to data handling are integral to the design of our classifier. Firstly, we apply the *k*-nearest neighbor (KNN) imputation technique to address structurally missing values—i.e., attributes that are consistently unreported by certain browser or mobile app configurations due to the inherent behaviour or policy of those environments, not accidental failures. Only numerical attributes are imputed, either originally continuous or transformed from high-cardinality categorical variables via frequency encoding. We use  $k = 5$  and the `nan_euclidean` distance metric, which computes distances over the available attributes only, making it well suited for partially missing fingerprints. Secondly, since XGBoost is primarily designed to work with numerical inputs,<sup>2</sup> we transform categorical variables following best practices: one-hot encoding for low-cardinality attributes and frequency encoding for high-cardinality ones. Frequency-encoded attributes yield continuous numerical values that reflect the statistical distribution of categories in the dataset, making them suitable for distance-based imputation methods.

Finally, we choose to develop two separate models, one for browsers and another for mobile apps. This decision is driven

<sup>2</sup>Note that the latest versions of XGBoost offer experimental support for categorical variables. However, this support is still limited.

Table I: Details of the ad campaigns used to populate our *fingerprint* datasets.

Campaign ID	Time Frame	Year	Ad Source	Type of Device	OS	CPM	delivered impressions	collected samples in our server	samples with advertising ID
adTag_adF-01	Feb 11-13	2022	webs; apps	mobile; desktop	all	0.02€	520007	381880	157293
adTag_adF-02	May 22-23	2022	webs; apps	mobile; desktop	Apple; Linux	0.03€	359909	302539	86471
adTag_adF-03	May 27-29	2022	webs; apps	desktop	macOS; Linux	0.32€	65066	53111	28294
adTag_adF-04	Jun 01-03	2022	webs; apps	desktop	macOS; Linux	0.09€	791092	590832	314750
adTag_adF-app_01	Jun 06-10	2022	apps	mobile	Android; iOS	0.03€	889335	594967	292382
adTag_adF-05	Jun 28-30	2022	webs; apps	desktop	macOS; Linux	0.38€	25892	23766	22241
adTag_adF-br_01	May 09-13	2023	webs	mobile; desktop	all	0.10€	983582	733229	327013
adTag_adF-br_02	Sep 19-24	2023	webs	mobile; desktop	all	0.06€	1134013	977378	469080
adTag_adF-app_02	Oct 03-07	2023	apps	mobile	Android; iOS	0.04€	634670	380042	261414

by the different number of attributes available in each case, 66 vs. 35. We also evaluate the option of building different independent models for individual device configurations in our datasets (e.g., a specific model for desktop devices running a Windows OS or for mobile Android apps), but we observe that the performance of *adF* is similar for the general and configuration-specific models, and thus we decide to implement the two referred general models. To further validate the robustness of these models, we compute additional evaluation metrics. The models achieve ROC AUC values of 90% and 88%, and balanced accuracy scores of 78% and 80% for browser-based and mobile app fingerprints, respectively.

## V. MEASUREMENTS AND DATASETS

To assess the vulnerability of different device configurations within browser-based environments to web fingerprinting, we deploy the *adF* system, as described earlier, in live ad campaigns. These campaigns are configured within the Sonata Platform, a Demand Side Platform (DSP) operated by TAPTAP Digital. Sonata is a mid-sized DSP that serves tens of millions of daily ads in 15 countries across Europe, Africa, and the Americas. The Data Collector retrieves the 66 attributes (highlighted in Table IV) from each of the ads delivered through these ad campaigns. Each sample is then fed into the Fingerprint Constructor and the Fingerprint Classifier to obtain the device’s fingerprint and determine whether it is unique or not. In addition, when available, we also collect the device’s advertising ID provided by the DSP as a feature.

The resulting dataset from these measurements, known as the *raw-browsers* dataset, comprises a total of 3.06M data samples collected at different time intervals (February-June 2022, May-June 2023, and September-October 2023). Each data sample includes the following fields:

1. device’s fingerprint obtained from the Fingerprint Constructor.
2. device’s advertising ID obtained from the DSP.
3. device’s type can take two values: mobile or desktop.
4. device’s OS captures the OS of the device: iOS or Android for mobile devices; Windows, macOS, or Linux for desktop devices.
5. browser represents the browser used by the device (e.g., Chrome, Safari, Firefox).

6. attributes is a data structure including the value of the 66 attributes.

We collect a second dataset, comprising 975k samples from mobile apps, distinct from the browser-based data. This dataset is referred to as the *raw-apps* dataset. The samples contain the same fields as previously described, with specific differences: (1) the field device’s type is always mobile in this case; (2) the field device’s OS can take only two values, iOS or Android; (3) the field browser identifies, in this case, the WebView of the Android or iOS system; (4) the data structure attributes includes information for 35 attributes.

Unfortunately, due to ethical and consistency concerns, we cannot utilize all the samples in the *raw-browsers* and *raw-apps* datasets. On the one hand, we aim to ensure our study does not subject users to risks they are not already exposed to. In our study, the risk primarily involves having a unique ID assigned. Consequently, we exclusively use data samples for which we possess an informed advertising ID, accounting for 45.88% and 56.80% in our *raw-browser* and *raw-apps* datasets, respectively. On the other hand, among these samples with an informed advertising ID, we only consider those for which the fingerprint value appears at least twice in the dataset. We do this because we cannot assess whether a fingerprint that appears just once is unique or not. The resulting datasets consist of 161k and 52k data samples, and we refer to them as the *fingerprint-browsers* and *fingerprint-apps* datasets, respectively. These datasets include two additional fields, which we term *ground-truth uniqueness* and *measured uniqueness*, respectively.

The *ground-truth uniqueness* is a binary variable that signifies whether a fingerprint is genuinely unique in our dataset (1) or not (0). Specifically, if all the samples in any of the *fingerprint* datasets sharing a specific *device’s fingerprint* value have the same advertising ID, the *ground-truth uniqueness* variable for these samples is set to 1 (indicating uniqueness), and it is set to 0 (indicating non-uniqueness) otherwise.

We employ the *ground-truth uniqueness* and the 66 attributes from the samples in the *fingerprint-browsers* dataset to train the XGBoost algorithm and derive the model to be implemented by our Fingerprint Classifier. The model’s output for each sample is then assigned to the *measured uniqueness* field of that sample within the *fingerprint-browsers* dataset.




Figure 2: x-axis:  $\mathcal{MV}$ ,  $\mathcal{TV}$  and  $\mathcal{A}$  performance of *adF* system for all the considered browser-based device and mobile app configurations; y-axis: number of samples for the browser and mobile app configurations considered for the analyses conducted in the paper.

For mobile apps, we follow the same procedure, using the 35 attributes from the *fingerprint-apps* dataset.

Finally, it is well-known that the market share of OSes and browsers is imbalanced. On the one hand, Chrome holds a two-thirds share of browser installations on both desktop and mobile devices. On the other hand, Windows OS dominates the desktop market with a share over 72,1% while the mobile device market is primarily split between Android (71%) and iOS (28,3%). To ensure that our dataset adequately represents the most popular browsers and OSes, we utilize the targeting capabilities provided by Sonata DSP to conduct specific ad campaigns tailored to particular OSes and browsers. Additionally, we run specific campaigns for mobile apps. The details of each campaign can be found in Table I. In addition, Figure 2 shows the number of samples in our *fingerprint-browsers* dataset across device configurations accounting for at least 725 samples. For statistical representativeness reasons, these are the configurations we consider in our analyses in the rest of the paper. In the case of mobile apps, our *fingerprint-apps* dataset includes 20,66% (10686) and 79,34% (41031) samples from iOS and Android apps, respectively.

## VI. VULNERABILITY & ACCURACY

In this section, we first define two metrics to assess the vulnerability of a given device configuration to web fingerprinting. In addition, we also introduce the accuracy metric used to assess the performance of *adF* system.

### A. Vulnerability Metrics

1) *Theoretical Vulnerability ( $\mathcal{TV}$ )*: We use *ground-truth uniqueness* from our *fingerprint* datasets to define the theoretical vulnerability ( $\mathcal{TV}$ ) of a given device configuration to web fingerprinting. This is defined as follows:

$$\mathcal{TV} = \frac{N_{tf}}{N_f} \quad (1)$$

Where  $N_f$  represents the total number of fingerprints and  $N_{tf}$  is the number of fingerprints with a *ground-truth uniqueness* equal to 1, for a given device configuration.  $\mathcal{TV}$  signifies the fraction of fingerprints that are unique in our dataset based on ground-truth data (this is why we use the term *theoretical* to refer to it).

2) *Measurable Vulnerability ( $\mathcal{MV}$ )*: It is analogous to  $\mathcal{TV}$ , but with a focus on *measured uniqueness* instead of *ground-truth uniqueness*. In particular, the equation that defines  $\mathcal{MV}$  is:

$$\mathcal{MV} = \frac{N_{mf}}{N_f} \quad (2)$$

Where  $N_f$  represents the total number of fingerprints and  $N_{mf}$  is the number of fingerprints with a *measured uniqueness* equal to 1, for a given device configuration.

Hence, while  $\mathcal{TV}$  captures the actual vulnerability of a given device configuration to web fingerprinting,  $\mathcal{MV}$  reflects the measurable vulnerability using our *adF* system, which is subject to errors (i.e., false positives and false negatives). If our system performs well, then  $\mathcal{TV}$  and  $\mathcal{MV}$  would exhibit similar values.

### B. adF system's accuracy

Our *adF* system essentially tackles a classification task, aiming to extract device fingerprints and determine their uniqueness. To assess performance, we use a standard classification metric, accuracy ( $\mathcal{A}$ ). Accuracy, in our scenario, represents the proportion of fingerprints correctly classified as unique (true positives) or non-unique (true negatives) by our system, relative to the total number of fingerprints ( $N_f$ ). A high accuracy suggests that even without access to advertising ID information, our system can effectively assess the vulnerability of different device configurations to web fingerprinting.

In addition, we can informally assess the performance of our web fingerprinting system by comparing the difference of the  $\mathcal{TV}$  and  $\mathcal{MV}$  values for different device configurations.

## VII. RESULTS

In this section, we first analyze the performance of *adF* system. Afterward, we analyze the vulnerability of prevalent device configurations formed by the combination of device type, operating system, and browser. Finally, we also analyze the vulnerability of mobile apps to web fingerprinting.

### A. Performance of our adF system

Figure 2 displays the values of  $\mathcal{TV}$ ,  $\mathcal{MV}$ , and  $\mathcal{A}$  for each configured mobile app and browser-based setup. *adF* consistently achieves high accuracy, with  $\mathcal{A}$  values exceeding 80% for all configurations. Moreover,  $\mathcal{MV}$  and  $\mathcal{TV}$  exhibit similar values across the configurations. Indeed, the maximum difference is only 11.02 percentage points for iOS apps. As such, we can confidently assert that *adF* is a reliable tool for assessing the vulnerability of different device configurations to web fingerprinting, even in the absence of ground-truth data like advertising IDs.

### B. Vulnerability of browser-based device configurations to web fingerprinting

The values of  $\mathcal{TV}$  and  $\mathcal{MV}$  presented in Figure 2 serve to assess the vulnerability of each of the browser-based configurations to web fingerprinting. Given the similarity between

$\mathcal{TV}$  and  $\mathcal{MV}$ , we discuss results based solely on  $\mathcal{MV}$  for the remainder of this section.

The first conclusion we extract from the results in Figure 2 is that mobile devices exhibit lower vulnerability to web fingerprinting when compared to desktop devices for a given browser (Chrome or Safari). For instance, in desktop devices running macOS, Safari's  $\mathcal{MV}$  is 5.5%, while in mobile devices running iOS, it is 3.4%. Likewise, Chrome offers lower  $\mathcal{MV}$  in mobile devices compared to desktops, regardless of the OS.

Second, the choice of operating system (OS) significantly influences defense against web fingerprinting. On mobile devices, regardless of the browser, iOS users enjoy better protection against web fingerprinting than their Android counterparts. For desktop devices, macOS and Linux outperform Windows in shielding against web fingerprinting. When considering the common browser, Chrome, across these OS categories,  $\mathcal{MV}$  rates are {95.7%, 85%, 78.1%} for {Windows, macOS, Linux}. However, unlike mobile devices, a single OS does not consistently offer superior protection regardless of the browser. For example, Chrome users on macOS and Edge users on Windows are more vulnerable compared to Firefox users on Linux.

Finally, when examining browsers individually, we find that Chrome ranks as the least secure option (except in the case of iOS). In contrast, Safari and Firefox provide notably superior protection against web fingerprinting. Note that we gather a few hundred samples from Firefox on both Windows and macOS, suggesting Firefox's strong performance also on these operating systems. To better understand the outstanding performance of Chrome on iOS, we refer the reader to our analysis in Section X. This result is due to the low cardinality and entropy of attributes accessible in Chrome for iOS devices, which does not occur in other OSes.

Putting the results into context:

The configurations in Figure 2 correspond to the most prevalent setups in the current web landscape, accounting for over 90% of all devices.

Our findings reveal significant variability in the vulnerability of these configurations to web fingerprinting. To put these results into perspective, it is essential to consider the market share, which is notably skewed, both in terms of OSes and browsers<sup>3</sup>.

In the mobile ecosystem, 71,01% and 28,30% of the devices run on Android and iOS, respectively. Regarding browsers, 64.37% of devices employ Chrome, and 24.98% opt for Safari. By combining OSes and browsers' market share, we identify three configurations that cover 95% of the mobile devices' market share: {Chrome, Android}  $\sim$ 67.7%, {Safari, iOS}  $\sim$ 25%, and {Chrome, iOS}  $\sim$ 3.3%. When we extrapolate the  $\mathcal{MV}$  values presented in Figure 2 to this market share, we observe that web fingerprinting is feasible for at least 40% of mobile devices when exposed to ads.

In the context of the desktop ecosystem, in terms of OSes, Windows, macOS, and Linux cover 72.09%, 16.85%, and

2.77% of the market share, respectively. When it comes to browsers, Chrome dominates with 65.59%, Edge follows with 10.66%, Safari has 10.60%, and Firefox stands at 6.95%. Therefore, the most prevalent desktop configurations are: {Chrome, Windows}  $\sim$ 54.5%, {Edge, Windows}  $\sim$ 10.7%, {Safari, macOS}  $\sim$ 10.6%, {Chrome, macOS}  $\sim$ 5.7%, {Chrome, Linux}  $\sim$ 1.2%, and {Firefox, Linux}  $\sim$ 1.2%. When extrapolating the results from Figure 2 to this specific market share, we find that the prevalence of the {Windows, Chrome} configuration leads to a situation where at least 66% of desktop devices can be effectively fingerprinted from ads.

### C. Mobile apps

Most mobile apps operate on top of Android or iOS WebView; therefore, they can be exploited to implement web fingerprinting. However, to the best of the authors' knowledge, there is no previous empirical study that confirms if this practice is viable. Figure 2 shows the value of  $\mathcal{TV}$ ,  $\mathcal{MV}$  and  $\mathcal{A}$  for mobile apps on *Android* vs. *iOS*.

First, we confirm that web fingerprinting is a viable practice in mobile applications. Furthermore, our experiments also confirm that *adF* is a valid tool to assess the vulnerability of mobile applications to web fingerprinting. In particular,  $\mathcal{A}$  is higher than 83% for both iOS and Android apps.

Second, Android apps are more exposed to web fingerprinting compared to iOS. In particular, when considering  $\mathcal{MV}$ , we find that 40.7% of Android users and 22% of iOS users can be uniquely fingerprinted using mobile apps' ads. If we extrapolate these results to the actual market share of Android and iOS (71.01% and 28.30%, respectively), we conclude that 35% of mobile devices can be uniquely identified using web fingerprinting techniques in mobile apps.

Finally, when comparing app versus browser vulnerability on mobile web, we find that while on Android, device vulnerability is significantly lower in mobile apps than browsers, on iOS, apps exhibit a notably higher vulnerability to web fingerprinting compared to browsers. This is a relevant finding for iOS mobile app developers who should consider when striving to enhance the security offered by their apps.

## VIII. METHODOLOGY LIMITATIONS

Our dataset is limited in device coverage, so our claim of fingerprint uniqueness is confined to our dataset's scope, and we cannot claim the global uniqueness of such fingerprints. This limitation stems primarily from our restricted budget for ad campaigns, amounting to 410,27 €. Also, ethical considerations led us to exclude almost half of the samples without an advertising ID in our dataset. Nevertheless, despite these constraints, our dataset's size (see Section VII) is comparable to previous works in the literature [8], [24]. It confirms that *adF* provides high accuracy, establishing it as a reliable system for evaluating device configuration vulnerability to web fingerprinting. Furthermore, our results highlight the significant variability in protection among different device configurations.

We also acknowledge that our methodology faces different issues that impede collecting the required data from every

<sup>3</sup>Market share statistics are based on data from <https://gs.statcounter.com/>, covering Feb 2022 to Oct 2023, providing a contemporary context to our dataset's time-frame.

single ad executing our Data Collector’s *adTag*. Some of these issues include: browsers running ad-blockers or disabling JavaScript, outdated browsers, limited browser API support, network problems in the connection between the ad and our server, etc. Moreover, the runtime required by our *adTag* represents another source of losses. Our *adTag* runs as long as the web-page (mobile app) where the ad is rendered is open on the user’s device. Hence, once the user closes the web-page (mobile app), the *adTag* execution is stopped. Callejo et al. [35] estimated that 75% of ads displayed on PCs and mobile devices remain active for at least 15 and 8 seconds, respectively. Although our *adTag* has an overall estimated runtime in the order of hundreds of milliseconds (depending on the browser and OS), we must also consider the rendering time of the website itself. All the mentioned causes lead to an average loss of 19% ad impressions from web-pages and of 33% ad impressions from mobile apps by comparing the campaign’s number of impressions reported by the DSP and the total samples collected in our back-end server (See Table I).

## IX. CAUSES OF WEB FINGERPRINTING VULNERABILITY

Different device configurations offer a significantly different vulnerability to web fingerprinting. In this section, we dissect the main factors that define the probability of a fingerprint being unique and, by extension, the resilience of a given device configuration to web fingerprinting.

### A. Reported attribute count ( $|N|$ )

An obvious factor that influences the resilience of a given device configuration (expressed through  $\mathcal{TV}$  or  $\mathcal{MV}$ ) is the length of the vector forming the fingerprint ( $|N|$ ) by such configuration to third-parties. For instance, in desktop devices,  $\{\text{Safari, macOS}\}$  provides information for only 39 out of the 66 attributes considered by *adF* whereas Firefox reports (regardless of the OS) data for 61 of these 66 attributes. Finally, Chrome provides, in any OS, information for all 66 attributes. Hence, in principle, Chrome users are more exposed to web fingerprinting since it is simpler to generate unique fingerprints using a larger pool of attributes.

Longer vectors generally result in a higher probability of acquiring unique fingerprints. In other words, device configurations that provide a greater number of attributes are typically more vulnerable to web fingerprinting. However, this is not the sole factor influencing the resilience to web fingerprinting. For instance, in the case of desktop configurations, if we use Safari’s anti-fingerprint settings in Chrome samples on Windows (i.e., we do not consider in Chrome those attributes not reported by Safari), the  $\mathcal{MV}$  ( $\mathcal{TV}$ ) for  $\{\text{Chrome, Windows}\}$  devices become 53,71% (60,35%). This represents an important improvement compared to the original vulnerability measured for  $\{\text{Chrome, Windows}\}$  configuration, 95,74% (95,81%). However, it is still far from  $\{\text{Safari, macOS}\}$  results ( $\mathcal{MV} = 5,48\%$ ,  $\mathcal{TV} = 8,36\%$ ). This example shows that two device configurations reporting exactly the same attributes present significantly different web fingerprinting vulnerabilities. Hence, there must be other factors influencing the web fingerprinting vulnerability.

### B. Attribute value range ( $|S|$ )

A second factor that defines the resilience of a device configuration to web fingerprinting is the *cardinality* of the reported attributes ( $|S|$ ), defined as the number of distinct values an attribute can assume. If attribute  $A$  can take 10 values whereas attribute  $B$  can take 20 values, then attribute  $B$  is (in principle) more useful to identify uniquely a user. For instance, in desktop devices, if we factor out those attributes taking just 1 or 2 values (i.e., deterministic or binary attributes), Chrome offers a larger cardinality for all attributes compared to Safari and Firefox. As an example, an attribute with a high cardinality like the *User-Agent* presents (731, 72, 82) potential values for (Chrome, Safari, Firefox), respectively.

However, although cardinality is crucial to determine an attribute’s capacity to fingerprint individual users, we find examples where it might be partially misleading. For example, the attribute *”Screen: pixel left”* can take potentially (2074, 286, 101) different values in (Chrome, Safari, Firefox) for desktop devices, respectively. In practice, (77,3%, 85,3%, 56,7%) of the samples in our *fingerprint-browsers* dataset take just a unique value (0) for this attribute in (Chrome, Safari, Firefox), respectively. Therefore, the discrimination power of an attribute to define unique fingerprints depends not only on its cardinality but also on the frequency of appearance of its potential values.

### C. Value frequency by attribute ( $h$ )

As proposed by Eckersley [5], we use the *entropy* to measure the frequency with which different values of a given attribute appear. The higher the entropy of an attribute, the more distributed the use of different values of such an attribute is.

We compute the Shannon entropy of the attributes considered in our study using the following expression:

$$H(X) = - \sum_{i=1}^n P(x_i) \cdot \log_b P(x_i) \quad (3)$$

Where  $H$  is the entropy,  $P$  is the probability distribution of the attribute’s values, and  $n$  is the total possible values for the attribute. The Shannon entropy is measured in bits and thus it uses a logarithm base equal to 2 ( $b = 2$ ).

To compare the entropy of attributes with different cardinality, we use a normalized version of Shannon’s entropy ( $h$ ) dividing the equation 3 by  $\log_2(M)$ , where  $M$  represents the number of values for each attribute.

$$h(X) = \frac{H(X)}{\log_2(M)} \quad (4)$$

$\log_2(M)$  is the maximum possible diversity index where the entropy is maximum, and all values of an attribute are equally frequent. Hence, the resulting normalized entropy ranges between 1 (the attribute’s largest discrimination power) and 0 (the attribute takes a unique value; thus, it is deterministic and useless for fingerprinting users).

However, the normalized entropy alone may again mislead us in identifying attributes with high discrimination power.

Table II: Discrimination power of attributes. For each attribute and device configuration, we report: (i) whether the attribute is reported or not, (ii) the cardinality ( $|S|$ ), (iii) the normalized entropy ( $h$ ). We highlight in bold those attributes whose reporting we propose to block in our solution ( $|S| > 25$  and a  $h \geq 0,10$ ).

Attributes	MOBILE								DESKTOP								APP							
	Android				iOS				Windows				macOS				Android		iOS					
	Chrome	MIUI	Chrome	Safari	Chrome	Edge	Chrome	Safari	Chrome	Firefox	Android	iOS												
$ S $	$h$	$ S $	$h$	$ S $	$h$	$ S $	$h$	$ S $	$h$	$ S $	$h$	$ S $	$h$	$ S $	$h$	$ S $	$h$							
UserAgent	6946	0,39	165	0,44	602	0,57	131	0,38	374	0,28	56	0,34	141	0,25	72	0,36	216	0,21	82	0,46	5050	0,73	106	0,33
<b>CPU cores</b>	8	0,02	2	0,01	4	0,03	4	0,05	24	0,14	9	0,22	13	0,14	4	0,12	24	0,13	9	0,23	6	0,02	2	0,03
<b>Device memory</b>	6	0,09	4	0,10	1	0,00	1	0,00	6	0,08	5	0,09	4	0,03	1	0,00	6	0,11	1	0,00	5	0,10	1	0,00
Screen: color depth	1	0,00	1	0,00	2	0,02	3	0,02	4	0,00	3	0,01	3	0,08	2	0,08	4	0,00	2	0,02	-	-	-	-
Screen: pixel left	4	0,00	1	0,00	1	0,00	2	0,00	1615	0,16	74	0,17	735	0,21	286	0,15	609	0,17	101	0,32	-	-	-	-
Screen: orientation angle	4	0,05	3	0,05	4	0,08	4	0,10	4	0,00	4	0,01	2	0,00	5	0,14	4	0,05	1	0,00	4	0,00	5	0,06
Screen: orientation type	4	0,05	3	0,05	4	0,08	4	0,10	4	0,00	4	0,00	3	0,00	5	0,14	4	0,07	1	0,00	-	-	-	-
Battery status: charging	3	0,04	3	0,05	1	0,00	1	0,00	3	0,03	3	0,05	3	0,08	1	0,00	3	0,07	1	0,00	-	-	-	-
Simultaneous touch points	4	0,02	1	0,00	3	0,00	4	0,00	19	0,02	6	0,05	2	0,00	3	0,08	8	0,09	2	0,00	2	0,03	2	0,01
<b>Media devices</b>	39	0,17	6	0,01	5	0,08	17	0,10	66	0,14	8	0,16	8	0,05	12	0,10	43	0,21	28	0,21	22	0,18	5	0,08
<b>Languages</b>	2748	0,31	22	0,13	27	0,15	39	0,18	1638	0,20	55	0,24	822	0,36	43	0,18	1355	0,39	40	0,25	285	0,17	42	0,20
PDF viewer enabled	3	0,01	2	0,05	3	0,08	3	0,08	3	0,04	3	0,01	3	0,01	3	0,09	3	0,08	3	0,09	-	-	-	-
<b>User Permissions state</b>	196	0,13	1	0,00	6	0,18	22	0,20	113	0,13	21	0,17	58	0,15	19	0,19	155	0,22	16	0,25	-	-	-	-
available height	456	0,32	41	0,31	30	0,29	32	0,25	549	0,26	79	0,42	582	0,52	310	0,45	590	0,44	206	0,63	227	0,33	20	0,24
available left	1	0,00	1	0,00	1	0,00	3	0,00	221	0,04	17	0,04	199	0,07	30	0,01	166	0,14	60	0,27	-	-	-	-
available top	1	0,00	1	0,00	1	0,00	3	0,00	484	0,03	13	0,02	181	0,18	13	0,11	170	0,12	38	0,28	-	-	-	-
available width	445	0,24	30	0,17	30	0,26	33	0,21	502	0,20	57	0,31	171	0,24	79	0,35	541	0,37	156	0,51	124	0,20	16	0,19
full screen enabled	2	0,06	2	0,10	3	0,01	1	0,00	3	0,05	2	0,10	2	0,07	3	0,11	3	0,07	3	0,10	2	0,06	1	0,00
<b>Storage: quota</b>	3650	0,60	198	0,67	12	0,03	18	0,03	23271	0,81	743	0,93	772	0,39	20	0,02	4664	0,74	91	0,21	1787	0,62	3	0,06
<b>navigator properties</b>	72	0,18	5	0,12	25	0,24	37	0,20	97	0,16	18	0,21	44	0,17	29	0,27	120	0,22	24	0,25	31	0,10	29	0,25
Plugins	1	0,00	1	0,00	2	0,08	2	0,08	51	0,05	4	0,01	8	0,02	21	0,17	47	0,09	4	0,09	-	-	-	-
Cookie enabled	2	0,00	2	0,01	2	0,08	2	0,07	2	0,00	2	0,01	2	0,00	2	0,07	2	0,00	2	0,07	-	-	-	-
MIME type	1	0,00	1	0,00	2	0,08	2	0,08	21	0,05	4	0,01	7	0,02	7	0,12	15	0,09	6	0,10	-	-	-	-
Time zone offset	20	0,15	7	0,12	13	0,13	19	0,15	24	0,08	9	0,17	15	0,15	14	0,11	14	0,09	9	0,14	15	0,13	12	0,18
<b>Canvas</b>	1360	0,46	70	0,45	90	0,31	311	0,35	619	0,33	134	0,37	247	0,39	236	0,47	1078	0,50	182	0,55	380	0,37	102	0,31
<b>Fonts</b>	14	0,05	2	0,06	13	0,10	17	0,11	8359	0,56	361	0,71	1934	0,52	78	0,24	1382	0,36	299	0,70	2	0,00	7	0,11
<b>Bluetooth availability</b>	3	0,06	2	0,05	1	0,00	1	0,00	3	0,08	3	0,10	3	0,08	1	0,00	3	0,10	1	0,00	-	-	-	-
<b>WebGL Extensions</b>	86	0,23	15	0,20	10	0,15	17	0,17	52	0,10	13	0,16	16	0,05	23	0,29	126	0,31	30	0,23	57	0,22	8	0,12
Audio formats: AAC	1	0,00	1	0,00	2	0,00	2	0,00	2	0,00	1	0,00	1	0,00	2	0,00	1	0,00	3	0,02	-	-	-	-
Audio formats: ACC	2	0,00	1	0,00	3	0,00	3	0,00	3	0,00	1	0,00	2	0,00	3	0,01	2	0,00	3	0,01	1	0,00	2	0,01
Audio cxt: base latency	116	0,22	9	0,22	8	0,06	11	0,05	18	0,02	3	0,03	9	0,09	13	0,11	76	0,20	5	0,04	70	0,25	7	0,05
Audio cxt: max channel count	1	0,00	1	0,00	4	0,02	4	0,03	5	0,02	5	0,02	14	0,01	7	0,03	5	0,00	5	0,02	-	-	-	-
Audio cxt: sample rate	4	0,01	2	0,00	6	0,05	7	0,04	10	0,03	5	0,04	9	0,08	10	0,09	7	0,07	3	0,09	2	0,01	6	0,04
Audio cxt: state	2	0,05	2	0,03	2	0,00	2	0,00	2	0,06	2	0,09	2	0,07	2	0,00	2	0,07	2	0,00	-	-	-	-
<b>WebGL (Rend - Param)</b>	525	0,46	76	0,52	10	0,09	18	0,11	2038	0,47	310	0,72	319	0,46	52	0,24	2360	0,65	137	0,54	234	0,40	7	0,05

For instance, an attribute with a normalized entropy close to 1 but a low cardinality, e.g., equal to 2, does not have a strong discrimination power. It just divides users into two groups. Therefore, the entropy alone is not sufficient to assess the discrimination power of an attribute.

#### D. Summary of the analysis

Our analysis above reveals that the discrimination power of an attribute is defined by three factors, it must: 1) be reported, 2) have high cardinality, and 3) have high entropy.

Table II provides a summary of these three factors ( $|N|$ ,  $|S|$ ,  $h$ ) for each attribute in the most relevant mobile app and browser-based configurations considered in this paper. Note that for attributes concatenated within a meta-attribute, we report the values for the meta-attribute.

The results in the table offer two main insights. On the one hand, we can assess that the configurations reported as most vulnerable to web fingerprinting in Section VII are characterized by a higher number of informed attributes with high discrimination power (high cardinality and normalized entropy). For instance, the Chrome, Windows configuration, identified as highly vulnerable, reports 35 out of the 35 attributes listed in Table II, with 8 of these attributes having a cardinality  $\geq 299$  and a normalized entropy  $\geq 0,20$ . Contrary, the Firefox, Linux configuration reports 30 attributes, with just 1 exceeding the indicated cardinality and normalized entropy thresholds. Likewise, it is evident that mobile configurations generally report fewer attributes than desktop configurations. Furthermore, attributes' cardinality and entropy tend to be lower in mobile configurations compared to desktop configurations.

On the other hand, Table II reveals those attributes with the most significant discrimination power (highlighted in bold). This represents very valuable information for browser and mobile apps developers, which can design solutions to mitigate the power discrimination of these attributes, reducing the vulnerability of the corresponding device configuration.

**Remark.** While this section focuses on static fingerprinting characteristics, Appendix B presents a longitudinal analysis of fingerprinting complexity over time, based on the joint entropy of full fingerprints across device configurations.

## X. COUNTERING WEB FINGERPRINTING

The efficiency of fighting fingerprint is nowadays (and will be) an added value feature, which might positively affect the market share of a given browser. The browser industry is aware of this. As described in Section II, they have (or are) developed (ing) different techniques to counter fingerprinting. These techniques, in essence, try to block the reporting of specific attributes or to reduce their cardinality and/or entropy. Our results in Section VII indicate that certain browsers already implement quite efficient solutions (e.g., Safari or Firefox) while others (e.g., Chrome) provide an almost negligible defense against web fingerprinting.

The question we raise is whether there is room for improvement. To answer this question, we want to avoid complex solutions, which, although probably very efficient, require major changes in browsers and mobile apps and thus are very unlikely to be adopted by developers. Following this principle, we then reformulate our question to whether there are simple (yet efficient) approaches that would substantially improve the




Figure 3: The effectiveness of the proposed anti-fingerprinting solution is indicated by  $\mathcal{MV}$ ,  $\mathcal{TV}$  and  $\mathcal{A}$  for each device configuration. Values above bars denote positive (blue) or negative (red) percentage variation compared to metrics without *ShieldF*.

resilience to web fingerprinting offered by current browsers. We anticipate that the answer is yes.

We propose *ShieldF*, a solution which relies on the most common technique already implemented by several browsers, this is, blocking/obfuscating the reporting of several attributes. This technique is simple, and browser developers are already using it, so it is likely that they are willing to adopt *ShieldF*. The key difference between *ShieldF* and existing solutions comes from the selection of the group of attributes to block. We propose to block those attributes which (1) have the largest discrimination power (i.e., highest cardinality and normalized entropy) and (2) do not affect (or have a very limited impact in) the end user experience<sup>4</sup>. To make an informed decision on the selected attributes to block, we rely on the results presented in Table II, which reports the discrimination power (cardinality and normalized entropy) of all considered attributes in our study. In particular, we propose to block those attributes with a cardinality  $> 25$ <sup>5</sup> and a normalized entropy  $\geq 0,10$ , which we are certain do not affect the browsing experience. The 12 selected attributes are highlighted in bold in Table II. Note this selection also implicitly accounts for correlated signals: we only consider an attribute mitigated when it cannot be inferred through other unblocked sources. Our proposal also applies to mobile apps with two exceptions: *User Permissions state* and *Bluetooth availability* attributes, because they are already unreported in mobile apps.

We assess the performance of *ShieldF* in a threefold manner. First, we measure the reduction in the web fingerprinting vulnerability it achieves for all considered device configurations. Second, we compare its performance with the one offered by the anti-fingerprinting solutions implemented by popular browsers (Chrome, Safari, Firefox). Note that we do not consider other browsers for different reasons: they have a low market share (MIUI, Tor, Brave), they are implemented using a Chromium Engine (Edge) and thus are expected to offer a similar performance to Chrome, their anti-fingerprinting solution affects the users' browsing experience (Tor), or their ads do not allow fingerprinting from ads (Brave). Furthermore, we

<sup>4</sup>For instance, an attribute revealing the screen size cannot be blocked even if its discrimination power is high. The screen size is required to adapt the content of a web-page to the screen size so it renders adequately.

<sup>5</sup>Except for *Bluetooth availability*, since it is characterized by a binary value.

also compare *ShieldF* performance with those solutions and techniques aimed at preventing unwanted web fingerprinting that adopt a similar methodology.

To further assess the impact of *ShieldF*, we compute the joint entropy of the blocked attributes before and after applying the countermeasure (see Appendix B for a formal definition). Results show a net drop of 2.15 bits in overall joint entropy (from 17.30 to 15.15), with an empirical  $p$ -value  $< 0.0001$  obtained via a permutation test under the null hypothesis of random assignment.

Note that for each specific solution considered in this section (those proposed by different browsers or ours), we re-train the Fingerprint Classifier from our *adF* fingerprinting system. Hence the reported  $\mathcal{MV}$  is the actual measurable vulnerability for each solution.

#### A. Fingerprinting vulnerability reduction

Figure 3 shows  $\mathcal{TV}$ ,  $\mathcal{MV}$  and  $\mathcal{A}$  for *ShieldF* for all the considered browser-based (mobile app) configurations. Moreover, for each configuration, the number reported in the figure over the  $\mathcal{TV}$  and  $\mathcal{MV}$ 's bar represents the reduction of the vulnerability (in percentage) compared to the original values of these metrics shown in Figure 2.

*ShieldF* reduces  $\mathcal{TV}$  and  $\mathcal{MV}$  in all considered configurations. In the case of  $\mathcal{TV}$ , such reduction ranges from 5% {Chrome, iOS} to 62% {Edge, Windows}.

#### B. Comparison to browsers' anti-fingerprinting proposals

We compare the performance of *ShieldF* with solutions currently implemented by the main browsers. We consider the {Desktop, Chrome, Windows} device configuration as the benchmark. We apply the most protective version of the solution (e.g., activating opt-in anti-fingerprint options) implemented by Safari, Firefox and Chrome introduced in Section II-C as well as our own solution to this configuration.

Figure 4 shows the values of  $\mathcal{TV}$ ,  $\mathcal{MV}$  and  $\mathcal{A}$  for each solution. The results indicate that Google's privacy sandbox and Firefox's solution lead to limited improvements. Safari's solution would improve the resilience to web fingerprinting in the {Desktop, Chrome, Windows} configuration by a reduction of  $\mathcal{MV}$  ( $\mathcal{TV}$ ) about 44% (37%), which is still far from the 60% (50%) improvement achieved by *ShieldF*.

#### C. Comparison with countermeasure

While there is no definitive approach that can prevent web fingerprinting while maintaining the richness of a browser, the scientific community has developed techniques and solutions aimed at mitigating the effects of web fingerprinting.

Some works describe and analyze a specific method, such as Font fingerprinting [17], Canvas fingerprinting [6], or WebGL fingerprinting [14]. Others take a broader view and quantify the combined use of functions and APIs most commonly employed for web fingerprinting [10]. While some solutions offer robust protection against browser fingerprinting, they often do so at the expense of usability [19]. The greatest challenge researchers face is providing complete coverage of




Figure 4: Comparing *ShieldF* to Safari, Firefox, and Google solutions. We assess  $\mathcal{M}$ ,  $\mathcal{T}$ , and  $\mathcal{A}$  on a benchmark setup,  $\{\text{Desktop, Windows, Chrome}\}$ .




Figure 5: Comparing *ShieldF* with techniques to mitigate web fingerprinting effects. We assess  $\mathcal{M}$ ,  $\mathcal{T}$ , and  $\mathcal{A}$  on a benchmark setup,  $\{\text{Desktop, Windows, Chrome}\}$ . Dashed line denotes  $\mathcal{M}$  of benchmark configuration  $\{\text{Desktop, Chrome, Windows}\}$ .

modified attributes, as even minor discrepancies can make users more visible to trackers [23].

Figure 5 displays the values of  $\mathcal{T}$ ,  $\mathcal{M}$ , and  $\mathcal{A}$  for the configuration  $\{\text{Desktop, Chrome, Windows}\}$  as a benchmark, as well as our solution, *ShieldF*, for this configuration. The results indicate that in the general context of fingerprinting, solutions targeting a specific method like Canvas Blocker [13], UniGL [14], or even FP-Random [18]—which focuses on defense techniques against Canvas fingerprinting, AudioContext fingerprinting, and unmasking browsers through JavaScript property ordering—offer protection that is seemingly marginal and makes almost no difference compared to not using them at all. Countermeasures like FP-Defense [38], FP-Block [39], FP-Defender [40], and DCB [41], which adopt a broader defense involving the combined use of JavaScript object-based functions and APIs, passive fingerprints, and HTTP headers, resulting in limited improvements in enhancing user privacy by preventing unwanted tracking. The FP-Defender solution would enhance resistance to web fingerprinting under the configuration of  $\{\text{Desktop, Chrome, Windows}\}$  by reducing  $\mathcal{M}$  ( $\mathcal{T}$ ) by approximately 11% (10%). However, this is still far from the improvement achieved by *ShieldF*, which is 60% (50%). On the other hand, the difference between  $\mathcal{M}$  and  $\mathcal{T}$  is larger with *ShieldF* compared to other solutions. This implies a lower accuracy of the *adF* fingerprinting system when our countermeasure *ShieldF* acts, as it has fewer attributes with high cardinality and entropy. This means that our *adF* fingerprinting system has a reduced capacity to identify unique fingerprints when *ShieldF* is active compared to other countermeasures.

One of the main challenges surrounding web fingerprinting and, thus, the development of defensive techniques is the continuous evolution of current web browser technologies. A proposed countermeasure can become obsolete in a matter of months, as each new browser version that modifies, adds, or even removes an object or API has a potential impact on the web fingerprinting ecosystem. For instance, the removal of access to NPAPI plugins brought the end of an important source of information for web fingerprinting and, consequently, for mitigation efforts of some countermeasures [17], [23], [39]. In contrast, the later introduction of features such as Canvas API [6], WebGL API [14], [40], Battery Status API [38], and Web Audio API [18], revealed new vulnerabilities in web fingerprinting exploitation. As a result, the web fingerprints of the past are different from those we see today and will likely differ from those we will encounter in the coming years. Section 11 discusses assessed countermeasures in more detail.

#### D. Empirical evaluation of *ShieldF*

We implement *ShieldF* as an extension for any Chromium-based browser (Google Chrome, Microsoft Edge, Opera, Brave, etc). First of all, we prove that the 12 selected attributes highlighted in Table II are effectively blocked by the *ShieldF* extension. Moreover, 15 beta-testers have installed the *ShieldF* extension and used it during at least a month. They have been instructed to report any issue arising from their browsing experience. At the time of submitting this paper, none of the beta-testers have reported any major issue. We acknowledge that this initial evaluation cannot be considered exhaustive due to the limited number of beta-testers, but it offers initial empirical evidence that *ShieldF* has a negligible impact in the browsing experience of users. The *ShieldF* extension is publicly available in the Chrome Web Store<sup>6</sup> to allow any interested users to test it.

-Summary: We conclude that a simple and easy-to-adopt solution by browsers and mobile app developers like *ShieldF* may substantially improve resilience against web fingerprinting. Note that our solution is based on a simple heuristic selection of attributes without the will of it being optimal, and thus, its performance can be considered a lower-bound improvement. For instance, each browser developer can take *ShieldF* code and the results reported in Table II as a starting point to define its optimized selection of attributes to be blocked in its browser.

Nevertheless, a crucial aspect of web fingerprinting and, consequently, mitigation solutions is the understanding that defense techniques quickly become deprecated. Therefore, this demands continuous monitoring of countermeasures to ensure they are regularly updated, guaranteeing that users remain protected against fingerprinting.

<sup>6</sup>*ShieldF* is available on: <https://chromewebstore.google.com/detail/shieldf/hdpnjiiibipldnnccgpcgilnpbdciia?hl=es&authuser=0>

<sup>7</sup> $H_m$  is the maximum possible diversity index where the entropy is maximum, and all values of an attribute are equally frequent.

Table III: Overview of large-scale fingerprinting studies: uniqueness ( $unq$ ), no. collected attributes ( $No. att$ ) and  $H_m$ .

	Eckersley [5]	Laperdrix et al. [7]		Gómez-Boix et al. [8]		Andriamilanto et al. [24]		adF		
	desktop	desktop	mobile	desktop	mobile	desktop	mobile	desktop	mobile	app
$unq$	94,20%	89,40%	81%	35,70%	18,50%	84,40%	39,90%	85,61%	49,93%	36,80%
No. att	6	17		17		262		66		35
$H_m^7$	18,843	16,859		20,98		21,983		16,64		15,66

## XI. RELATED WORK

### A. Fingerprinting on the web

The profound privacy risks tied to web fingerprinting have motivated the research community to address the topic.

The existing literature has extensively examined the feasibility of employing various JavaScript APIs for fingerprinting devices such as Canvas [6], WebGL API [9], available fonts [8], Web Audio API [9], Battery Status API [42] or extensions [33], as well as combinations of attributes; Mayer [4] conducted, back in 2009, one of the pioneering works to uniquely identify individual devices. In particular, the author collected the content from *navigator*, *screen*, *navigator.plugins*, and *navigator.mimeTypes* attributes to compute the devices fingerprints. In addition, the research community has conducted different studies based on the measurement and web crawling techniques to assess to what extent web fingerprinting is used [6].

Most of the previous studies focus on desktop devices. However, the rapid growth of the mobile web and the development of its underlying technology has enabled the possibility to implement fingerprinting techniques to identify mobile devices uniquely [43]. Eubank et al. [44] conducted one of the earliest studies on mobile web tracking, analyzing cookies and HTTP headers used by browsers and revealing distinctions between the fingerprints of mobile and desktop devices.

An overview of this literature can be found in the surveys by Laperdrix et al. [30] and Zhang et al. [45]).

In the context of our work, few papers have attempted large-scale analysis to assess users' exposure to web fingerprinting. A seminal work in the field by Eckersley [5] performed the first large-scale analysis of web fingerprinting. By leveraging the Panopticlick [5] website, the authors analyzed attributes collected from HTTP headers, JavaScript APIs, and Flash and Java plugins. More recently, Laperdrix et al. [7] performed the first large-scale analysis considering fingerprints from both mobile and desktop devices. They based their study on data collected from the AmIUnique website [7]. These aforementioned works followed a similar approach, involving websites where volunteer users connect. These websites run scripts to compute the device's fingerprint and assess the uniqueness of each. Instead, Gómez-Boix et al. [8] and Andriamilanto et al. [24] collected fingerprints through a script deployed on one of the top 15 French websites, as per the Alexa traffic rank. Similar to [8], [24], Li et al. [25] deployed a fingerprinting tool on a real-world website and studied the effectiveness of web fingerprinting by collecting millions of samples.

We present a novel approach for measuring web fingerprinting using ads as measurement venues rather than websites. This offers some advantages such as: 1) ability to define

targeted measurement experiments for specific device configurations, which allows for first time to reveal the vulnerability at the granularity of individual device configurations; 2) independence from the willingness of volunteers or website owners to provide data; 3) capacity to audit the vulnerability to web fingerprinting in mobile apps for the first time.

Table III compares the major large-scale studies on web fingerprinting, including our own, across three dimensions: uniqueness, defined as the fraction of unique identified fingerprints for both desktop and mobile devices; number of collected attributes; and maximum diversity index. We observe that our methodology aligns well with the findings of the most recent study [24] in terms of uniqueness.

### B. Fingerprinting countermeasures

The research community has been devoted to finding solutions to countermeasure the fingerprint. Finding a definitive method or solution that avoids fingerprinting does not exist without compromising the user's browsing experience. A widespread idea among the research community is to approach this issue to seek collective immunity, hide the user in the larger set of users, and try to standardize their attributes by overriding, randomizing, spoofing, etc. We find in the literature two types of proposals to counter fingerprinting.

The first group proposes different types of manipulation of the reported attributes to third-parties [10], [11]. The solutions proposed by browser developers and introduced in Section II fall into this group. We also find proposals from the research community. For instance, Blink [10] or Mimic [12] propose to modify the content of the fingerprint by spoofing the value of certain attributes, whereas Canvas Blocker [13] blocks the whole Canvas. UniGL [14] proposes to universalize the fingerprint by removing discrepancies for the WebGL API between devices. The approach we propose in this paper, *ShieldF*, belongs to this group. Inspired by the work of Laperdrix et al. [7], we do a thorough exploration of the attributes reported by browsers and mobile apps to third-parties to block those with a larger discrimination power, granted that they do not affect the end user's browsing experience.

The second group aims at identifying and blocking fingerprinting scripts. For instance, FP-Inspector [15] uses machine learning classification algorithms and applies a layered set of restrictions to detect fingerprinting scripts. Privacy Badger [16] prevents the execution of known fingerprinting scripts in the browser. Also, a broader extension such as NoScript [19], which blocks the execution of unwanted JavaScript scripts, can be applied in the context of fingerprinting. This group of solutions are complementary to the one presented in this paper.

## XII. CONCLUSION

In this paper, we present *adF*, a novel system to assess the vulnerability of different device configurations to web fingerprinting. *adF* leverages the widespread use of online ads as measurement venues for collecting device attributes from code embedded in ads. This gathered information is then processed to identify unique fingerprints. We use *adF* in several real ad campaigns, enabling us to estimate the current vulnerability of devices to web fingerprinting. The results obtained are concerning. In particular, 66% of desktop devices are vulnerable to web fingerprinting. In the case of mobile devices, 40% are vulnerable to web fingerprinting via browser-rendered ads, and 35% through mobile apps.

Furthermore, we conduct an in-depth analysis of the discrimination power associated with different attributes reported by browsers and mobile apps. These findings hold significant value for developers of browsers and mobile apps seeking to implement more effective countermeasures against web fingerprinting by blocking or controlling those attributes with a higher discrimination power. Our detailed analysis also allows us to define a straightforward yet efficient solution to web fingerprinting. This solution proposes to block those attributes with the largest discrimination power, whose blocking does not affect the user experience in browsers or mobile apps. Our analysis shows that our proposal (easily adopted by browsers and mobile app developers) outperforms all the anti-fingerprinting solutions offered by major browser (Chrome, Safari and Firefox), reducing device vulnerability to web fingerprinting up to 62%. Given that collecting attributes from scripts embedded in ads may present more limitations compared to scripts embedded in web-pages' HTML code, we believe that the results presented in this paper can be considered a lower bound of the vulnerability of devices to traditional web fingerprinting.

## ACKNOWLEDGMENT

This work is financially backed through Grant PRITIA-CLOUD from the UNICO-CLOUD program, financed by the Ministry for Digital Transformation and Civil Service and the European Union NextGenerationEU/PRTR. Further funding is allocated to the the framework of the Recovery, Transformation and Resilience Plan funds, financed by the European Union (Next Generation) through the grant ANálisis y miTigación de riesgos de seguridad y privaCIad asociados a la explotación de datos PersonAles en OTTs (ANTICIPA). Additional support is provided by the framework of the Recovery, Transformation and Resilience Plan, funded by the European Union (Next Generation), under Grant *Cátedras ENIA 2022 para la creación de cátedras universidad-empresa en IA*, specifically AImpulsa: Cátedra UC3M - Universia de Economía del Dato y la Inteligencia Artificial Responsable aplicada a la Creación Exponencial de Valor.

## REFERENCES

- [1] "Alphabet annual report 2023," <https://abc.xyz/assets/43/22/5deeff4fbec54014ae97b340c22/34ac6dab5f586b2e6e008b99fe683e35.pdf>, Feb 2024.
- [2] "Meta reports fourth quarter and full year 2023 results; initiates quarterly dividend," <https://investor.fb.com/investor-news/press-release-details/2024/Meta-Reports-Fourth-Quarter-and-Full-Year-2023-Results-Initiates-Quarterly-Dividend/default.aspx>, Feb 2024.
- [3] T. Bujlow, V. Carela-Español, J. Sole-Pareta, and P. Barlet-Ros, "A survey on web tracking: Mechanisms, implications, and defenses," *Proceedings of the IEEE*, vol. 105, no. 8, pp. 1476–1510, 2017.
- [4] J. R. Mayer, "Any person... a pamphleteer": Internet anonymity in the age of web 2.0," *Undergraduate Senior Thesis, Princeton University*, vol. 85, 2009.
- [5] P. Eckersley, "How unique is your web browser?" *Privacy Enhancing Technologies*, p. 1–18, 2010.
- [6] 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.
- [7] P. Laperdrix, W. Rudametkin, and B. Baudry, "Beauty and the beast: Diverting modern web browsers to build unique browser fingerprints," in *2016 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2016, pp. 878–894.
- [8] A. Gómez-Boix, P. Laperdrix, and B. Baudry, "Hiding in the crowd: an analysis of the effectiveness of browser fingerprinting at large scale," in *Proceedings of the 2018 world wide web conference*, 2018, pp. 309–318.
- [9] Y. Cao, S. Li, E. Wijmans *et al.*, "(cross-) browser fingerprinting via os and hardware level features," in *NDSS*, 2017.
- [10] P. Laperdrix, W. Rudametkin, and B. Baudry, "Mitigating browser fingerprint tracking: multi-level reconfiguration and diversification," in *2015 IEEE/ACM 10th International Symposium on Software Engineering for Adaptive and Self-Managing Systems*. IEEE, 2015, pp. 98–108. [Online]. Available: <https://hal.inria.fr/hal-01121108/document>
- [11] N. M. Al-Fannah and C. Mitchell, "Too little too late: can we control browser fingerprinting?" *Journal of intellectual capital.*, vol. 21, no. 2, pp. 165–180, 2020.
- [12] B. A. Azad, O. Starov, P. Laperdrix, and N. Nikiforakis, "Taming the shape shifter: Detecting anti-fingerprinting browsers," in *DIMVA 2020-17th Conference on Detection of Intrusions and Malware & Vulnerability Assessment*, 2020.
- [13] "Firefox extension to block the canvas api," <https://addons.mozilla.org/fr/firefox/addon/canvasblocker/>.
- [14] S. Wu, S. Li, Y. Cao, and N. Wang, "Rendered private: Making {GLSL} execution uniform to prevent {WebGL-based} browser fingerprinting," in *28th USENIX Security Symposium (USENIX Security 19)*, 2019, pp. 1645–1660. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity19/presentation/wu>
- [15] 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.
- [16] "Privacy badger official website," <https://privacybadger.org/>.
- [17] N. Nikiforakis, W. Joosen, and B. Livshits, "Privariator: Deceiving fingerprinters with little white lies," in *Proceedings of the 24th International Conference on World Wide Web*, 2015, pp. 820–830. [Online]. Available: <https://dl.acm.org/doi/pdf/10.1145/2736277.2741090>
- [18] P. Laperdrix, B. Baudry, and V. Mishra, "Fprandom: Randomizing core browser objects to break advanced device fingerprinting techniques," in *International Symposium on Engineering Secure Software and Systems*. Springer, 2017, pp. 97–114. [Online]. Available: <https://hal.inria.fr/hal-01527580/document>
- [19] "What is it? - noscript: Own your browser!" <https://noscript.net/>.
- [20] D. Cameron, "Apple declares war on 'browser fingerprinting,' the sneaky tactic that tracks you in incognito mode," <https://gizmodo.com/apple-declares-war-on-browser-fingerprinting-the-sneak-1826549108/>, Jun 2018.
- [21] D. Camp, "Firefox now available with enhanced tracking protection by default plus updates to facebook container, firefox monitor and lockwise," <https://blog.mozilla.org/blog/2019/06/04/firefox-now-available-with-enhanced-tracking-protection-by-default/>.
- [22] A. W. Writer, "User-agent reduction," <https://developer.chrome.com/docs/privacy-sandbox/user-agent/>.
- [23] K. Boda, Á. M. Földes, G. G. Gulyás, and S. Imre, "User tracking on the web via cross-browser fingerprinting," in *Nordic conference on secure it systems*. Springer, 2011, pp. 31–46. [Online]. Available: <https://link.springer.com/content/pdf/10.1007%2F978-3-642-29615-4.pdf>
- [24] N. Andriamilanto, T. Allard, G. Le Guelvouit, and A. Garel, "A large-scale empirical analysis of browser fingerprints properties for web

- authentication,” *ACM Trans. Web*, vol. 16, no. 1, sep 2021. [Online]. Available: <https://doi.org/10.1145/3478026>
- [25] S. Li and Y. Cao, “Who touched my browser fingerprint? a large-scale measurement study and classification of fingerprint dynamics,” in *Proceedings of the ACM Internet Measurement Conference*, ser. IMC '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 370–385. [Online]. Available: <https://doi.org/10.1145/3419394.3423614>
- [26] M. Gabryel, K. Grzanek, and Y. Hayashi, “Browser fingerprint coding methods increasing the effectiveness of user identification in the web traffic,” *Journal of Artificial Intelligence and Soft Computing Research*, vol. 10, 2020.
- [27] Maksim, “Google ending third-party cookies in chrome,” <https://www.cookiebot.com/en/google-third-party-cookies/>, 2022.
- [28] “Advertising id - play console help,” <https://support.google.com/googleplay/android-developer/answer/6048248?hl=en>.
- [29] Brave, “Fingerprinting protection mode · brave/browser-laptop wiki,” <https://github.com/brave/browser-laptop/wiki/Fingerprinting-Protection-Mode/>.
- [30] P. Laperdrix, N. Bielova, B. Baudry, and G. Avoine, “Browser fingerprinting: A survey,” *ACM Trans. Web*, vol. 14, no. 2, apr 2020. [Online]. Available: <https://doi.org/10.1145/3386040>
- [31] S. Khattak, D. Fifield, S. Afroz, M. Javed, S. Sundaresan, V. Paxson, S. J. Murdoch, and D. McCoy, “Do you see what i see? differential treatment of anonymous users,” in *NDSS Symposium'16 - San Diego, CA, US*. Internet Society, 2016. [Online]. Available: <https://discovery.ucl.ac.uk/id/eprint/1472066/>
- [32] Y. Mobile, “Privacy changes in app marketing for 2024,” <https://mobilemarketingmagazine.com/privacy-changes-in-app-marketing-for-2024/>, December 2023, accessed: 2024-03-21.
- [33] P. Laperdrix, O. Starov, Q. Chen, A. Kapravelos, and N. Nikiforakis, “Fingerprinting in style: Detecting browser extensions via injected style sheets,” in *30th USENIX Security Symposium (USENIX Security 21)*, 2021, pp. 2507–2524.
- [34] P. Callejo, A. Pastor, R. Cuevas, and Á. Cuevas, “Q-tag: A transparent solution to measure ads viewability rate in online advertising campaigns,” in *Proceedings of the 15th International Conference on Emerging Networking Experiments And Technologies*, 2019, pp. 151–157.
- [35] P. Callejo, C. Kelton, N. Vallina-Rodriguez, R. Cuevas, O. Gasser, C. Kreibich, F. Wohlfart, and Á. Cuevas, “Opportunities and challenges of ad-based measurements from the edge of the network,” in *Proceedings of the 16th ACM Workshop on Hot Topics in Networks*, 2017, pp. 87–93.
- [36] I. Media, “Study: Online ads need to be viewable for 14 seconds to be seen,” <https://www.inskinmedia.com/post/study-online-ads-need-to-be-viewable-for-14-seconds-to-be-seen/>, Oct 2020.
- [37] S. Arlot and A. Celisse, “A survey of cross-validation procedures for model selection,” 2010.
- [38] V. L. Ajay and A. M. Guptha, “A defense against javascript object-based fingerprinting and passive fingerprinting,” in *2022 International Conference on Computing, Communication, Security and Intelligent Systems (IC3SIS)*. IEEE, 2022, pp. 1–6.
- [39] C. F. Torres, H. Jonker, and S. Mauw, “Fp-block: usable web privacy by controlling browser fingerprinting,” in *European Symposium on Research in Computer Security*. Springer, 2015, pp. 3–19. [Online]. Available: [https://link.springer.com/chapter/10.1007/978-3-319-24177-7\\_1](https://link.springer.com/chapter/10.1007/978-3-319-24177-7_1)
- [40] D. Moad, V. Sihag, G. Choudhary, D. G. Duguma, and I. You, “Fingerprint defender: defense against browser-based user tracking,” in *International Symposium on Mobile Internet Security*. Springer, 2021, pp. 236–247.
- [41] P. Baumann, S. Katzenbeisser, M. Stopczynski, and E. Tews, “Disguised chromium browser: Robust browser, flash and canvas fingerprinting protection,” in *Proceedings of the 2016 ACM on Workshop on Privacy in the Electronic Society*, 2016, pp. 37–46.
- [42] C. Diaz, L. Olejnik, G. Acar, and C. Castelluccia, “The leaking battery: a privacy analysis of the html5 battery status api,” *Lecture notes in computer science*, vol. 9481, pp. 254–263, 2015.
- [43] A. Tiwari, J. Prakash, A. Rahimov, and C. Hammer, “Understanding the impact of fingerprinting in android hybrid apps,” in *2023 IEEE/ACM 10th International Conference on Mobile Software Engineering and Systems (MOBILESoft)*, May 2023, pp. 28–39.
- [44] C. Eubank, M. Melara, D. Perez-Botero, and A. Narayanan, “Shining the floodlights on mobile web tracking—a privacy survey,” in *Proceedings of the IEEE Workshop on Web*, vol. 2, 2013.

- [45] D. Zhang, J. Zhang, Y. Bu, B. Chen, C. Sun, T. Wang *et al.*, “A survey of browser fingerprint research and application,” *Wireless Communications and Mobile Computing*, vol. 2022, 2022.

## BIOGRAPHY



projects.

**Miguel A Bermejo-Agueda** received the B.Sc. degree in Telecommunications Engineering from the Universidad Politécnica de Madrid, and the M.Sc. degree in Data Science from the Universitat Oberta de Catalunya. He is currently pursuing a Ph.D. degree in Telematics Engineering at the Universidad Carlos III de Madrid. He is an author of conference papers, such as ACM WWW, PoPETs. His areas of interest include Internet measurements, online advertising, web transparency, and privacy. As a Ph.D. researcher, he has participated in EU H2020



She has also participated in EU H2020 projects. Her areas of interest include Internet measurements, online advertising, and web transparency.

**Patricia Callejo** is a Visiting Professor at Universidad Carlos III de Madrid. She obtained her M.Sc. (2016) and Ph.D. (2020) from Universidad Carlos III de Madrid in the field of Telematics Engineering. She was granted by RIPE Academic Cooperation Initiative (RACI) on RIPE 76, held in Marseille, France in 2018. The same year, she completed an internship at the International Computer Science Institute (ICSI) at UC Berkeley (USA) as part of her Ph.D. Dr. Callejo is author of conference papers such as ACM HotNets, ACM CoNEXT, and WWW.



Management at Aalborg University in 2006. He is co-author of more than 70 papers in prestigious international journals and conferences, including ACM CoNEXT, WWW, Usenix Security, ACM HotNets, IEEE Infocom, ACM CHI, IEEE/ACM TON, IEEE TPDS, CACM, PNAS, Nature Scientific Reports, PlosONE or Communications of the ACM. His main research interests encompass online advertising, web transparency, personalization and privacy, online social networks, and Internet measurements.

**Rubén Cuevas** is an Associate Professor at Universidad Carlos III de Madrid and serves as the Deputy Director of the UC3M-Santander Big Data Institute. From January to December 2012, he was Courtesy Assistant Professor in the Computer and Information Science department at University of Oregon. He obtained his Ph.D. and M.Sc. in Telematics Engineering and M.Sc. in Telecommunications Engineering at Universidad Carlos III de Madrid in 2010, 2007 and 2005 respectively. Furthermore, he received his M.Sc. in Network Planning and



WWW, ACM CONEXT, ACM CHI and USENIX Security. His research interests focus on Internet measurements, web transparency, privacy, and P2P networks. Dr. Cuevas is recipient of the Best Paper Award in ACM MSWiM 2010 and the 2018 Emilio Aced Research Prize granted by the Spanish Data Protection Agency.

**Ángel Cuevas** is an Associate Professor at Universidad Carlos III de Madrid. He received the B.Sc. degree in Telecommunication Engineering, and the M.Sc. and the Ph.D. degrees in Telematics Engineering from the Universidad Carlos III de Madrid, in 2006, 2007, and 2011, respectively. He is a co-author of more than 70 papers in prestigious international journals and conferences, such as the CACM, PNAS, IEEE/ACM Transactions on Networking, ACM Transactions on Sensor Networks, IEEE Network, IEEE Communications Magazine,



**Álvaro Mayol** currently serves as the Managing Director and Partner at TAPTAP Digital, having previously played a crucial role as Chief Product Officer (CPO), leading the Product and Engineering Department. Before holding these leadership positions, Álvaro served as the Chief Technology Officer (CTO) at the same company. He is a Telecommunications Engineer educated at the Universidad Politécnica de Madrid and the Illinois Institute of Technology. He began his professional career at Indra Systems as a systems engineer in the Electronic

Warfare and Military Communications department, where he participated in notable projects such as the A400M and ESSOR. Later, he led the Solutions and Engineering department at Geometry Global, a WPP agency. In 2014, TAPTAP Digital acquired the company that Álvaro co-founded, which specialized in proximity marketing and analytics based on iBeacon and other wireless technologies.

## APPENDIX A ETHICS CONSIDERATIONS

The research and experiments described in this paper obtained the IRB approval of our institution through the Data Protection Officer, a member of the Ethics Committee responsible for approving projects and experiments with potential data protection implications.

Note that we take specific measures to prevent new risks for users (see Section V). In particular, we only consider fingerprints for those users for which we have an advertising ID. By doing so, we avoid creating any unique ID for users who do not have one. Moreover, we do not process any data from devices that have the do-not-track flag active in their browser.

Finally, our *adTag* sends a total of 38.7kB to our backend server. This is a negligible amount of data, which has a minimal impact on the overall data consumption of the device.

## APPENDIX B TEMPORAL STABILITY OF CONFIGURATION-LEVEL WEB FINGERPRINTING

To complement our core evaluation and address potential concerns regarding the temporal robustness of fingerprinting, we conduct a longitudinal analysis across four timeframes in our dataset: 02 - 2022, 05/06 - 2022, 05 - 2023, and 09/10 - 2023. Unlike prior works as Laperdrix et al. (2020) that examine the persistence of individual fingerprints over sessions, our approach focuses on the temporal evolution of the fingerprinting landscape at the configuration level, defined as the tuple  $\{\text{browser}, \text{OS}\}$ , as observed through repeated ad delivery campaigns.

While Section IX.C analyzes the entropy of individual attributes, here we extend the perspective using the *joint entropy* of all fingerprinting attributes to capture the overall structural complexity of the fingerprint space over time. Formally, the Shannon joint entropy of a set of attributes  $X_1, X_2, \dots, X_n$  is:

$$H(X_1, X_2, \dots, X_n) = - \sum_{(x_1, x_2, \dots, x_n) \in \mathcal{X}} P(x_1, x_2, \dots, x_n) \times \log_2 P(x_1, x_2, \dots, x_n) \quad (5)$$




Figure 6: Joint entropy ( $\mathcal{H}$ ) over time for selected browser–OS configurations meeting sample stability criteria. Desktop setups show consistently higher entropy, while iOS-based configurations maintain lower and stable values.

where  $\mathcal{X}$  is the set of all possible combinations of attribute values, and  $P(x_1, x_2, \dots, x_n)$  is the joint empirical probability of each such combination. This measure captures the richness and dispersion of fingerprints independently of class imbalance or sample size.

To ensure reliability, we include only configurations that (i) maintain at least 400 samples per timeframe, and (ii) exhibit a sample count variability (standard deviation over mean) below 70% across timeframes. These constraints reduce bias from underrepresented or unstable datasets. The resulting set includes configurations for iOS and MacOS, which meet these criteria consistently. Figure 6 summarizes their temporal evolution.

Among desktop setups,  $\{\text{Chrome}, \text{MacOS}\}$  maintains high joint entropy (11.3–12.2 bits), indicating a persistent and diverse fingerprint space—suggesting little evolution in defenses. In contrast,  $\{\text{Safari}, \text{MacOS}\}$  remains more stable and less expressive (10.0–11.0 bits), consistent with reduced API exposure.

On mobile,  $\{\text{Safari}, \text{iOS}\}$  shows a notably stable and low joint entropy (9.9–11.8 bits), reaffirming its well-known fingerprinting resistance.  $\{\text{Chrome}, \text{iOS}\}$  exhibits the lowest entropy values across all configurations (<10.3 bits), likely due to iOS sandboxing that limits Chrome’s attribute exposure. This behavior contrasts with desktop counterparts and supports previous findings that attribute diversity on iOS is intentionally constrained by platform-level controls.

No major non-monotonic shifts are observed among the selected configurations. Minor fluctuations (e.g., a dip for  $\{\text{Chrome}, \text{iOS}\}$  in 05 - 2023) correspond to lower sample sizes and should be interpreted with caution. These trends align with the findings of Laperdrix et al. (2020), who report that fingerprint changes are often driven by changes in browser APIs, software updates, and hardware context—not necessarily by user behaviour or population shifts.

In summary, fingerprinting vulnerability does evolve over time, albeit unevenly across platforms and browsers. Joint

entropy emerges as a robust metric to capture these temporal dynamics and validates the ability of *adF* to operate as a longitudinal and configuration-aware fingerprinting auditor under real deployment conditions.

#### APPENDIX C

##### ATTRIBUTES CONSIDERED BY THE *adF* SYSTEM

This appendix includes Table IV providing detailed information of all the attributes initially considered to build the *adF* system.

