

Talon: An Automated Framework for Cross-Device Tracking Detection

Konstantinos Solomos
FORTH, Greece
solomos@ics.forth.gr

Sotiris Ioannidis
FORTH, Greece
sotiris@ics.forth.gr

Panagiotis Ilija
Univ. of Illinois at Chicago, USA
pilia@uic.edu

Nicolas Kourtellis
Telefonica Reasearch, Spain
nicolas.kourtellis@telefonica.com

Abstract

Although digital advertising fuels much of today's free Web, it typically does so at the cost of online users' privacy, due to the continuous tracking and leakage of users' personal data. In search for new ways to optimize the effectiveness of ads, advertisers have introduced new advanced paradigms such as cross-device tracking (CDT), to monitor users' browsing on multiple devices and screens, and deliver (re)targeted ads in the most appropriate screen. Unfortunately, this practice leads to greater privacy concerns for the end-user.

Going beyond the state-of-the-art, we propose a novel methodology for detecting CDT and measuring the factors affecting its performance, in a repeatable and systematic way. This new methodology is based on emulating realistic browsing activity of end-users, from different devices, and thus triggering and detecting cross-device targeted ads. We design and build Talon¹, a CDT measurement framework that implements our methodology and allows experimentation with multiple parallel devices, experimental setups and settings. By employing Talon, we perform several critical experiments, and we are able to not only detect and measure CDT with average AUC score of 0.78-0.96, but also to provide significant insights about the behavior of CDT entities and the impact on users' privacy. In the hands of privacy researchers, policy makers and end-users, Talon can be an invaluable tool for raising awareness and increasing transparency on tracking practices used by the ad-ecosystem.

1 Introduction

Online advertising has become a driving force of the economy, with digital ad spending already surpassing the spending for TV-based advertising in 2017 [32], and expected to reach \$327 billion in 2019 [42]. This is because online advertising can be easily tailored to, and target specific audiences. In order to personalize ads, advertisers employ various tracking practices to collect user behavioral and browsing data.

¹<https://en.wikipedia.org/wiki/Talos>




Figure 1: High level representation of cross-device tracking.

Until recently, the tracking of a user was confined to the physical boundary of each one of her devices. However, as users typically own multiple devices [2, 54], advertisers have started employing advanced targeting practices specifically designed to track and target users across **all** their devices. These efforts indicate a radical shift of the ad-targeting paradigm, from *device-centric* to *user-centric*. In this new paradigm, an advertiser tries to identify which devices (e.g., smartphone, tablet, laptop) belong to the same user, and then target her across all devices with ads related to her overall online behavior. Figure 1 illustrates a typical cross-device tracking (CDT) scenario, where a user is targeted with relevant ads in her second device (desktop), due to the behavior exhibited to the ad-ecosystem from her first device (mobile).

A recent FTC Staff Report [51] states that CDT can be deterministic or probabilistic, and companies engaging in such practices typically use a mixture of both techniques. Deterministic tracking utilizes 1st-party login services that require user authentication (e.g., Facebook, Twitter, Gmail). These 1st-party services often share information (e.g., a unique identifier) with 3rd-parties, enabling them to perform a more effective CDT. In the case of probabilistic CDT, there are no shared identifiers between the users' devices, and 3rd-parties attempt to identify which devices belong to the same user by considering network access data, common behavioral patterns in browsing history, etc. In fact, to understand the degree to which CDT trackers appear on the Web, we

measured their frequency of appearance on Alexa Top-10k websites: companies performing probabilistic CDT can be found in $\sim 27\%$ of the websites, and when also considering deterministic CDT, this coverage reaches $\sim 80\%$. Also, several advertising companies such as Criteo [16], Tapad [53], Drawbridge [20] etc., claim that they can track users across devices with very high accuracy (e.g., Drawbridge’s Cross-Device connected consumer graph is 97.3% accurate [19]).

In spite of its big impact on user privacy, apart from some empirical evidence about CDT, there is only a limited work investigating it. In the most close work to ours, Zimmeck et al. [56], designed an algorithm that correlates mobile and desktop devices into pairs by considering devices’ browsing history and IP addresses. While this approach shows that correlation of devices is possible when such data are available, it does not provide an approach for detecting and measuring CDT. In fact, to the best of our knowledge, there is no existing approach to audit the probabilistic CDT ecosystem and the factors that impact its performance on the Web. *Our work is the first to propose a novel methodology that enables auditing the CDT ecosystem in an automated and systematic way. In effect, our work takes the first and crucial step in understanding the inner workings of the CDT mechanics and measure different parameters that affect how it performs.*

The methodology proposed in this work is based on the following idea: we want to detect when CDT trackers successfully correlate a user’s devices, by identifying cross-device targeted behavioral ads they send, i.e., ads that are delivered on one device, but have been triggered because of the user’s browsing on a different device. In order to design this methodology, we first study browsing data of real users with multiple devices from [56] and extract topics of interest and other user behavioral patterns. Then, to make trackers correlate the different devices of the end-user and serve cross-device targeted ads, we employ artificially created personas with specific interests, to emulate realistic browsing activity across the user devices as extracted from the real data.

We build Talon, a novel framework that materializes our methodology in order to collect, categorize and analyze all the ads delivered to the different user devices, and evaluate with simple and advanced statistical methods the potential existence of CDT. Through a variety of experiments we are able to measure CDT with an average AUC of 0.78-0.96. Specifically, in the simplest experiment, where the user exhibits significant browsing activity mainly from the mobile device, the average value of AUC is 0.78 for the 10 different behavioral profiles used. When the user exhibits significant browsing activity from both devices (mobile and desktop), with a matching behavioral profile, we observe CDT with an average AUC of 0.83. In the case of visiting specifically chosen websites that employ multiple known CDT trackers, we achieve AUC score of 0.96. We also find that browsing in incognito can reduce the effect of CDT, but does not eliminate it, as trackers can perform device matching based

only on the current browsing session of the user, and not all her browsing history. Finally, we compare the data collected with our real user-driven artificial personas (such as CDT trackers found, types of ads detected, etc.) with corresponding distributions observed in the real user data from [56], offering a strong validation to the realistic design of Talon.

Overall, our main contributions in this work are:

- Design a novel, real data-driven methodology for detecting CDT by triggering behavioral cross-device targeted ads on one user device, according to specifically-crafted emulated personas, and then detecting those ads when delivered on a different device of the same user.
- Implement Talon, a practical framework for CDT measurements. Talon has been designed to provide scalability for fast deployment of multiple parallel device instances, to support various experimental setups, and to be easily extensible.
- Conduct a set of experiments for measuring the potential existence of CDT in different types of emulated users, with an average AUC score of 0.78-0.96, and investigate the various factors that affect its performance under different classes of experimental setups and configurations.

2 Related Work

The ad-industry continuously develops new mechanisms for making ads more relevant and effective. Such mechanisms include the delivery of contextual, targeted-behavioral, and retargeted ads. However, in order to serve such highly related ads, advertisers often employ questionable and privacy intrusive techniques for collecting user information. They typically apply techniques for tracking user visits across different websites, which allow them to reconstruct parts of the users’ browsing history. To that end, numerous works [48, 34, 39, 43, 17, 46, 45, 47, 44] investigate the various approaches employed by trackers, and propose protection mechanisms. Also, a large body of work investigates targeted behavioral advertising with regards to different levels of personalization, based on the type of information used to target the user [9, 5, 55], and its effectiveness [25, 35, 26, 12, 7, 29].

Some studies investigate CDT that utilizes technologies such as ultrasound and Bluetooth, and measure the prevalence of such approaches [40, 6, 33]. A study by Brookman et al. [10] provided initial insights about the prevalence of CDT on the web, identified 3rd-party CDT trackers and examined the transparency of the employed techniques.

Zimmeck et al. [56] conducted a small-scale exploratory study on CDT based on the observation of cross-device targeted ads in two paired devices using the same IP address (mobile and desktop) over the course of two months. Following this exploration, they collected real users’ browsing histories and device information and designed an algorithm

that correlates the devices into pairs. This approach shows that network information and browsing history can be used for correlating user devices, and thus potentially for CDT.

In general, research around CDT is still very limited; in fact, only [56, 10] initially studied some of its aspects, but without proving its actual existence or providing a methodology for detecting and measuring it. Our work builds on these early studies on CDT, as well as past studies on web tracking, and proposes a methodology that enables systematic investigation and measurements for detecting probabilistic CDT.

3 A methodology to measure CDT

The proposed methodology emulates realistic browsing activity of end-users across different devices, and collects and categorizes all ads delivered to these devices based on the intensity of the targeting. Finally, it compares these ads with baseline browsing activity to establish if CDT is present or not, at what level, and for which types of user interests.

3.1 Design Principle

In general, the CDT performed by the ad-ecosystem is a very complex process, with multiple parties involved, and a non-trivial task to dissect and understand. To infer its internal mechanics, we rely on probing the ecosystem with consistent and repeatable inputs (I), under specific experimental settings (V), allowing the ecosystem to process and use this input via transformations and modeling (F), and produce outputs we can measure on the receiving end (Y):

$$(I, V) \xrightarrow{F} Y$$

In this expression, the unknown F is the probabilistic modeling performed by CDT entities, allowing them to track users across their devices. Following this design principle, our methodology allows to push realistic input signals to the ad-ecosystem via website visits, and measure the ecosystem's output through the delivered ads, to demonstrate if F enabled the ecosystem to perform probabilistic CDT. An overview of our methodology is illustrated in Figure 2.

3.2 Design Overview

3.2.1 Input Signal (I)

To trigger CDT, we first need to inject to the ad-ecosystem some activity from a user's browsing behavior (I). This *input* can be visits (i) to pages of interest (e.g., travel, shopping), or (ii) to control pages of null interest (e.g., weather pages). Intuitively, the former can be used first to demonstrate particular behavior of a user from a given device (mobile), and the latter afterwards for collecting ads delivered as the *output* of the ecosystem (Y) due to I, to that device, or other device of the same user (desktop).

Persona Pages. We extract real users' interests from the dataset provided by Zimmeck et al. [56] and leverage an approach similar to Carrascosa et al. [12] to emulate browsing behavior according to specific web categories, and create multiple, carefully-crafted *personas* of different granularities. This design makes the methodology systematic and repeatable and produces realistic browsing traffic from scripted browsers. For each persona, our approach identifies a set of websites (dubbed as *persona pages*) that have, at the given time, active ad-campaigns. This "*training activity*" aims to drive CDT trackers into possible *device-pairing* between the user's two devices with high degree of confidence.

Control Pages. Following past works [12, 7], all devices in the system collect ads by visiting neutral websites that typically serve ads not related to their content, thus, reducing bias from possible behavioral ads delivered to specific type of websites. We refer to these websites as *control pages*. We detail the design of personas and control pages in § 4.1.

3.2.2 Experimental Setup (V)

No 1st-party logins. Since we focus on probabilistic CDT, we assume that the emulated user does not visit or log into any 1st-party service that employs deterministic CDT and thus, there is no common identifier (e.g., email address, social network ID) shared between the user's devices.

Devices, IP addresses & Activity. The approach we follow is based on triggering and identifying behavioral cross-device targeted ads, and specifically ads that appear on one of the user's devices, but have been triggered by the user's activity on a different device. For this trigger to be facilitated, the ad-ecosystem must be provided with hints that these two devices belong to the same user. Zimmeck et al. [56] suggest that in many cases, the devices' IP address is adequate for matching devices that belong to the same user. Also, according to relevant industrial teams [38, 4] more signals can be used, such as location, browsing, etc., for device matching.

Following these observations, our methodology requires a minimum of three different devices: one mobile device and two desktop computers, with two different public IP addresses. We assume that two devices (i.e., the mobile and one desktop) belong to the same user, and are connected to the same network. That is, these devices have the same public IP address, are active in the same geolocation as in a typical home network, and will be considered by the ad-ecosystem as producing traffic from the same user. The second desktop (i.e., *baseline PC*), which has a different IP address, is used for receiving a different flow of ads while replicating the browsing of the user's desktop (i.e., *paired PC*). This control instance is used for establishing a baseline set of ads to compare with the ads received by the user's paired PC.

CDT Direction. In principle, the design allows the investigation of both directions of CDT. That is, users may first browse on the mobile device, and then move to their desk-

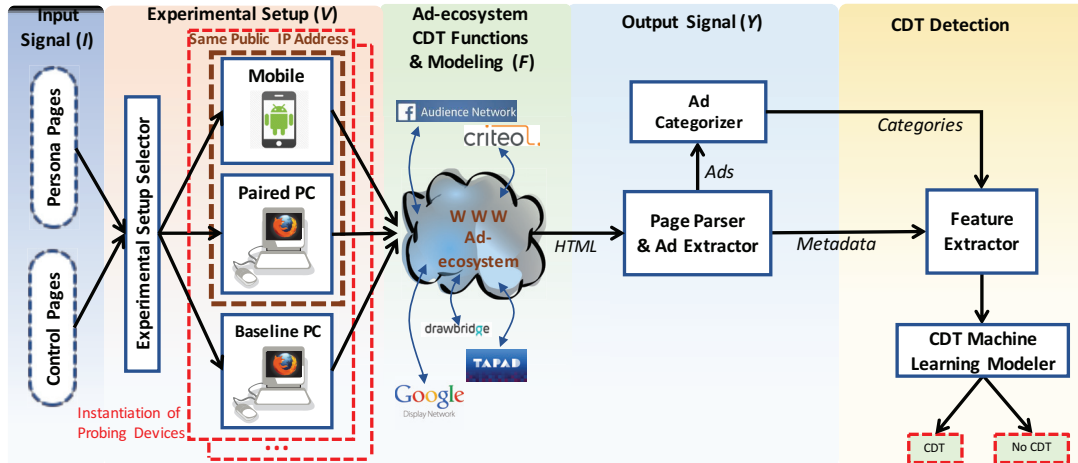


Figure 2: High level representation of methodology design principles and units for CDT measurements.

top, and vice versa. However, since ad-targeting companies such as AdBrain and Criteo support that the direction from mobile to desktop is more suitable for cross-device retargeting [49, 3, 15], in this work we focus on the mobile to desktop direction (*Mob* \rightarrow *PC*). In essence, the mobile device performs a specifically instructed web browsing session to establish the persona, by visiting the set of *persona pages*, i.e., *training* phase; then, the two desktop computers perform web browsing, i.e., *testing* phase, where they visit the set of *control pages* and collect the delivered ads. The browsing performed by the desktops is synchronized by means of visiting the same pages and performing the exact same clicks.

3.2.3 Output Signal (*Y*)

In order to handle the Output Signal and transform it appropriately, we design and implement two different components: (i) Page Parser & Ad Extractor and (ii) Ad Categorizer. The first is responsible for the identification and extraction of ad elements inside the webpages. The module uses string matching techniques and a public list of common ad-domains (*Easylist* [21]) to identify the delivered ads. The second module assigns a keyword on each ad identified on the previous step, based on its type and content (e.g., “Online Shopping”, “Fashion”, “Recreation”, etc.). Using both modules, we store the ads delivered in all devices of our experimental setup along with their categories, as well as data related to the activity of the devices that attracted these ads.

3.2.4 CDT Detection

Comparing Signals. Various statistical methods can be used to associate the input signal *I* of persona browsing in the mobile device, with the output signal *Y* of ads delivered to the potentially paired-PC. For example, simple methods that perform similarity computation between the two signals in a

given dimensionality (e.g., Jaccard, Cosine) can be applied. These methods, as well as typical statistical techniques (e.g., permutation tests) capture only one dimension of each input/output signal and thus, might not be suitable for measuring with confidence the high complexity of the CDT signal. In this case, more advanced methods can be employed, such as Machine Learning techniques (ML) for classification of the signals as similar enough to match, or not. In our analysis, we mainly focus on ML to compute the likelihood of the two signals being the product of CDT, as it takes into consideration this multidimensionality in the feature space. We describe the modeling and methods used for ML in § 4.4.

4 Framework Implementation

A high level overview of our methodology, and its materialization by our framework Talon, is presented in Figure 2 and described in § 3. In the following, we provide more details about its building blocks, and argue for various design decisions taken while implementing this methodology into the fully-fledged automated system.

4.1 Input Signal: Control Pages & Personas

Persona Pages. A critical part of our methodology is the design and automatic building of realistic user personas. Each persona has a unique collection of visiting links, that form the set of *persona pages*. Since we do not know in advance which e-commerce sites are conducting cross-device ad-campaigns, we design a process to dynamically detect active persona pages of given interest categories. Our approach for persona generation is shown in Figure 3.

We first use the list of topics of Zimmeck et al. [56], that describe real user’s online interests. We perform a clustering based on the content of each interest and label the clusters appropriately (e.g., we group together: “Shopping” and

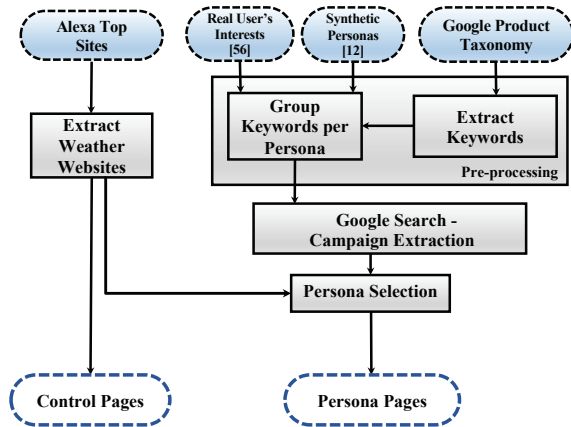


Figure 3: Persona design and automatic generation.

“Beauty and Fashion” under the label: “Shopping and Fashion”). Then, we use the persona categorization of Carrascosa et al. [12] for their top 50 personas, and select only those personas that describe similar interests with the previously formed list. For the resulting intersection of personas from the two lists, we iterate through the Google Product Taxonomy list [27] to obtain the related keywords for each one.

For increasing the probability to capture active ad-campaigns that can potentially deliver ads to the devices, we use Google Search as it reveals campaigns associated with products currently being advertised. That is, if a user searches for specific keywords (e.g., “men watches”), Google will display a set of results, including sponsored links for sites conducting campaigns for the terms searched. In this way, we use the keywords set for each persona, as extracted above, and transform them into search queries by appending common string patterns such as “buy”, “sell”, and “offers”. This process is repeated until between five and ten unique domains per persona are collected. If the procedure fails, no persona is formed.

As the effectiveness of a persona depends on the active ad-campaigns at the given time, in our experiments, we deploy personas in 10 categories related to shopping, traveling, etc. (full list shown in Table 3 in Appendix). With this procedure, we manage to design personas similar enough with real users, as well as with emulated users designed in previous works [12, 7, 8, 56].

Control Pages. For retrieving the delivered ads (after any type of browsing), we employ a set of webpages that contain: (i) easily identifiable ad-elements and (ii) a sufficient number of ads that remains consistent through time. These pages have neutral context and do not affect the behavioral profile of the device visiting them. For most of the experiments in § 5, we use a set of five popular weather websites² as control pages, similarly to [12]. We manually confirmed the neutral-

²accuweather.com, wunderground.com, weather.com, weather-forecast.com, metcheck.com

ity of these pages, by observing no contextual ads delivered to them. When visiting the set of control pages, our methods extract and categorize all the ads received, in order to identify those that have been potentially resulted from CDT.

4.2 Experimental System Setup

The experimental setup contains different types of units, connected together for replicating browsing activity on multiple devices. Typically, CDT is applied on two or more devices that belong to the same user, such as a desktop and a mobile device. Thus, the system contains emulated instances of both types, controlled by a number of experimental parameters.

Devices & Automation. The desktop devices are built on top of the web measurement framework OpenWPM [22]. This platform enables launching instances of the Firefox browser, performs realistic browsing with scrolling, sleeps and clicks, and collects a wide range of measurements in every browsing session. It is also capable of storing the browser’s data (cookies, local cache, temporary files) and exports a browser profile after the end of a browsing session, which can be loaded in a future session. With these options, we can perform *stateful* experiments, as a typical user’s web browser that stores all the data through time, or *stateless* experiments to emulate browsing in incognito mode.

For the mobile device, we use the official Android Emulator [28], as well as the Appium UI Automator [50] for the automation of browsing. We build the mobile browsing module on top of these components to automate visits to pages via the Browser Application. This browsing module provides functionalities for realistic interaction with a website, e.g., scrolling, click and sleep rate. Similarly to the desktop, it can run either in a *stateful* or *stateless* mode.

Experimental Setup Selector. As shortly described in § 3, we need two phases of browsing to different types of webpages (*training* and *testing*), in order to successfully measure CDT. For that reason, we set the two browsing phases in the following way: During the training phase, the selected device visits the set of *Persona Pages* for a specific duration, referred to as training time (t_{train}). The test phase is the set of visits to *control pages* for the purpose of collecting ads. During this phase, we control the duration of browsing (i.e., t_{test}). The experimental setup selector controls various parameters such as: which type of device will be trained and tested, the times t_{train} and t_{test} , the sequence of time slots for training and testing from the selected device, number of repetitions of this procedure, etc.

Timeline of phases. Each class of experiments is executed multiple times (or runs), through parallel instantiations of the user devices within the framework (as shown in Figure 2). Each experimental run is executed following a timeline of phases as illustrated in Figure 4. This timeline contains N sessions with three primary stages in each: Before, Mobile, and After. The *Before* (B_i) stage is when the two desktop de-




Figure 4: Timeline of phases for CDT measurement.
 M_i : mobile training time t_{train} + testing time t_{test} ;
 $B_i(A_i)$: desktop testing time t_{test} before (after) mobile phase;
 W : wait time (t_{wait}); R : rest time (t_{rest}); t_{s_i} : time of session i .

vices perform a parallel test browsing, with a duration of t_{test} time, to establish the state of ads before the mobile device injects signal into the ad-ecosystem. The *Mobile* (M_i) stage is when the mobile device performs a training browsing for t_{train} time, and a test browsing for t_{test} time. This phase injects the signal from the mobile during training with a persona, but also performs a subsequent test with control pages to establish the state of ads after the training. Finally, the *After* (A_i) stage is when the two desktops perform the final test browsing, with the same duration t_{test} as in *Before* (B_i) stage, to establish the state of ads after the mobile training.

After extensive experimentation, we found that a minimum training time $t_{train}=15$ minutes and testing time $t_{test}=20$ minutes are sufficient for injecting a clear signal over noise, from the trained device to the ad-ecosystem. There is also a waiting time ($t_{wait}=10$ minutes) and resting time ($t_{rest}=5$ minutes) between the stages of each session, to allow alignment of instantiations of devices running in parallel during each session. In total, each session lasts 1.5 hours and is repeated $N=15$ times during a run. Through the experimental setup selector, we define the values of such variables (t_{train} , t_{test} , t_{wait} , t_{rest} , N , type of device), offering the researcher the flexibility to experiment in different cases of CDT.

4.3 Output Signal

Page Parser. This component is activated when the visited page is fully loaded and no further changes occur on the content. To collect the display ads, we first need to identify specific DOM elements inside the visited webpages. This task is challenging due to the dynamic Javascript execution and the complex DOM structures generated in most webpages. For the reliable extraction of ad-elements and identification of the landing pages,³ we follow a methodology similar to the one proposed in [36]. The functionality of this component is to parse the rendered webpage and extract the attributes of display ads, which also contain the landing pages.

Ad Extractor. In most modern websites, the displayed ads are embedded in *iFrame* tags that create deep nesting layers, containing numerous and different types of elements. How-

ever, the ads served by the control pages are found directly inside the *iFrames* so the module does not have to handle such complex behavior. Therefore, the module firstly identifies all the active *iFrame* elements and filters out the invalid ones that have either empty content or zero dimensions. Then, it retrieves the *href* attributes of image and flash ads and parses the URLs, while searching for specific string patterns such as *adurl=*, *redirect=*, etc. These patterns are typically used by the ad-networks for encoding URLs in webpages. Next, the module forms the list of candidate landing pages, which are then processed and analyzed to create the set of true landing pages. The Ad Extractor is fully compatible with the crawlers, and does not need to perform any clicks on the ad-elements, since it extracts only the landing pages' URLs directly from the rendered webpage. After collecting the candidate landing pages, the module filters them with the *EasyList* [21], similarly to previous works [7, 22], and stores only the true active ad-domains. Finally, the Page Parser & Ad Extractor module also stores *metadata* from the crawls such as: time and date of execution, number of identified ads, number of categories, type and phase of crawl, etc.

Ad Categorizer. To associate landing pages or browsing URLs with web categories, we employ the McAfee Trusted-Sources database [41], which provides URLs organized into categories. This system was able to categorize 96% of the landing pages of our collection into a total of 76 unique categories, by providing up to four semantic categories for each page, while the remaining 4% domains were manually classified to the categories above. The final output contains the landing pages of collected ads, along with their categories.

4.4 CDT Detection

Probabilistic CDT is a kind of task generally suitable for investigation through ML. Previous work [56] and industry directions [38, 4] claim that probabilistic device-pairing is based on specific, well-defined signals such as: IP address, geolocation, type and frequency of browsing activity. Since we control these parameters in our methodology, by definition we construct the ground truth with our experimental setups. That is, we control (i) the devices used, which are potentially paired under a given IP address, geolocation and browsing patterns, (ii) the control instance of baseline desktop device, and (iii) the browsing with the personas.

Before applying any statistical method, every instance of the input data has to be transformed into a vector of values; each position in the vector corresponds to a feature. Features are different properties of the collected data: browsing activity of a user during training time, experimental setup used (persona, etc.), time-related details of the experiment, as well as information about the collected ads, which is the output signal received from the given browsing activity. These features can be studied systematically to identify statistical association between the input and output signals, given an ex-

³Destination websites the user is redirected to when clicking on the ads.

perimental setup. In effect, our feature space is comprised of a union of these vectors, since all features are either controlled, or measurable by us (detailed description of the features is given in Appendix, Table 5). The only unknown is whether the ad-ecosystem has successfully associated the devices, and if it has exhibited this in the output signal via ads.

One Dimension Statistical Analysis. At the first level of analysis, to measure the similarity of distribution of ads delivered in the different devices, we compare the signals using a two-tailed permutation test and reject the null hypothesis that the frequency of ads delivered (for a given category) comes from the same distribution, if the t-test statistic leads to a p-value smaller than a significance level $\alpha < 0.05$.

Multidimensional Statistical Analysis. Given that a uni-dimensional test such as the previous one does not take into account the various other features available in each experiment, we further consider ML, which take into account multidimensional data, to decide if the ads delivered in each device are from the same distribution or not. We transform the problem of identifying if the previously exported vectors are similar enough, into a typical binary classification problem, where the predicted class describes the existence of pairing or not, that may have occurred between the mobile device and one of the two desktop devices. As a paired combination we consider the desktop device that exists under the same IP address with the mobile device. The “not paired” combination is the mobile device and the baseline desktop. The analysis is based on three classification algorithms with different dependences on the data distributions. An easily applied classifier that is typically used for performance comparison with other models, is the Gaussian Naive Bayes classifier. Logistic Regression is a well-behaved classification algorithm that can be trained, as long as the classes are linearly separable. It is also robust to noise and can avoid overfitting by tuning its regularization/penalty parameters. Random Forest and Extra-Trees classifiers, construct a multitude of decision trees and output the class that is the mode of the classes of the individual trees. Also they use the Gini index metric to compute the importance of features.

A fundamental point when considering the performance evaluation of ML algorithms is the selection of the appropriate metrics. Pure Accuracy can be used, but it’s not representative for our analysis, since we want to report the most accurate estimation for the number of predicted paired devices, while at the same time measure the absolute number of miss-classified samples overall. For this reason, metrics like Precision, Recall and F_1 -score, and the Area Under Curve of the Receiver Operating Curve (*AUC*) are typically used, since they can quantify this type of information.

5 Experimental Evaluation

We use the Talon framework to perform various experiments and construct different datasets for each. Since every ex-

Table 1: Characteristics of the datasets used in each setup (S) of experiments. $S=\{1,2,3\}$ are the setups of experiments in § 5.2, § 5.3 and § 5.4, respectively; t_{total} : the total duration of experiment; t_{train} : the training duration; t_{test} : the testing duration; I: independent personas; C: data combined from personas; SF: stateful browser; SL: stateless browser; B: boosted CDT browsing.

S	Personas	Runs	t_{train}	t_{test}	t_{total}	Samples	Features
1a	10 (I, SF)	4	15min	20min	37 days	240	1100
1b	10 (C, SF)	-	-	-	-	2400	2201
2a	2 (I, SF)	4	480min	30min	6 days	192	600
2b	2 (C, SF)	-	-	-	-	384	750
2c	2 (I, SF, B)	4	480min	30min	6 days	192	500
2d	2 (C, SF, B)	-	-	-	-	384	576
3a	5 (I, SL)	2	15min	20min	9 days	120	450
3b	5 (C, SL)	-	-	-	-	600	880

perimental setup has different experimental parameters (i.e., training and testing time, number of personas, browsing functionalities), the datasets vary in terms of samples size and feature space. The datasets collected during our experiments and used in our analysis are presented in the Table 1.

5.1 Does IP-sharing allow CDT?

A first set of preliminary experiments were performed to demonstrate that our platform can (i) successfully identify and collect the ads delivered to our multiple devices (mobile and desktops), (ii) inject browsing signal from a device, thus biasing it to have a realistic persona and (iii) lead to matching/pairing of devices, which could be due to same behavioral ads, retargeting ads or CDT.

First, we use a simple experimental setup: we connect three instances of desktop devices and one mobile device under the same IP address. We create one persona (as in § 4.1), with an interest in “Online Shopping-Fashion, Beauty”, and following the described timeline of phases, we run this experiment for two days. Then, we perform one-dimensional statistical analysis, as introduced in § 4.4, and find that there is no similarity between the mobile with any of desktop devices (null hypothesis rejected with highest p-value=0.030), while all desktop distributions are similar to each other (null hypothesis accepted with lowest p-value=0.33). These statistical results indicate that there is no clear device-pairing (at the level of ad distribution for the given persona), and that we should consider controlling more factors to instigate it.

Consequently, we expand this experiment by also training one of the desktop devices using the same persona as with mobile. By repeating the same statistical tests, we find that the mobile and desktop with the same browsing behavior receive ads coming from the same distribution (null hypothesis accepted with lowest p-value=0.84), while the other desktop devices show no similarity with each other or the mobile (null hypothesis rejected with highest p-value=0.008). This

result indicates that browsing behavior under a shared IP address can boost the signal towards advertisers, which they can use to apply advanced targeting, either as CDT, or retargeting on each device or a mixture of both techniques.

Finally, these preliminary experiments and statistical tests provide us with evidence regarding the effectiveness of our framework to inject enough browsing signal from different devices under selected personas. Our framework is also able to collect ads delivered between devices, that can be later analyzed and linked back to the personas. Those are fundamental components for our system and importantly they are potentially causing CDT between the devices involved. Next, we present more elaborate experimentations with our framework, in order to study CDT in action.

5.2 Does short-time browsing allow CDT?

Independent Personas: Setup 1a. This experimental setup emulates the behavior of a user that browses frequently about some topics, but in short-lived sessions in her devices. Given that most users do not frequently delete their local browsing state, this setup assumes that the user’s browser stores all state, i.e., cookies, cache, browsing history. This enables trackers to identify users more easily across their devices, as they have historical information about them. In this setup, every experimental run starts with a clean browser profile; cookies and temporary browser files are stored for the whole duration of the experimental run (stateful). We use all personas of Table 3, and the data collection for each lasts 4 days.

We perform the same statistical analysis as in § 5.1, and find that in 4/10 personas, the mobile and paired desktop ads are similar (null hypothesis accepted with lowest p-value=0.13), while the mobile and baseline desktop ad distributions are different (null hypothesis is rejected with highest p-value=0.009). This inconsistency is reasonable since the statistical analysis is based only on one dimension (the frequency count of types of ads appearing in the devices), which may not be enough for fully capturing the existence of device-pairing. For this reason, we choose to use more advanced, multidimensional ML methods which take into account the various variables available, to effectively compare the potential CDT signals received by the two devices.

The classification results of the Random Forest (best performing) algorithm are reported in Table 2. We use AUC score as the main metric in our analysis, since the ad-industry seems to prefer higher Precision scores over Recall, as the False Positives have greater impact on the effectiveness of ad-campaigns.⁴ As shown in Table 2, the model achieves high AUC scores for most of the personas, with a maximum value of 0.84. Specifically, the personas 2, 4 and 8 scored

⁴Tapad [1] mentions: “Maintaining a low false positive rate while also having a low false negative rate and scale is optimal. This combination is a strong indicator that the Device Graph in question was neither artificially augmented nor scrubbed.”

Table 2: Performance evaluation for Random Forest in Setups 1a and 1b. Left value in each column is the score for Class 0 (C0=*not paired desktop*); right value for Class 1 (C1=*paired desktop*).

Persona (Setup)	Precision		Recall		F ₁ -Score		AUC
	C0	C1	C0	C1	C0	C1	
1 (1a)	0.89	0.60	0.57	0.90	0.70	0.72	0.73
2 (1a)	0.84	0.78	0.81	0.82	0.82	0.80	0.82
3 (1a)	0.81	0.73	0.78	0.76	0.79	0.74	0.76
4 (1a)	0.87	0.78	0.87	0.78	0.87	0.78	0.82
5 (1a)	0.94	0.65	0.68	0.93	0.79	0.76	0.80
6 (1a)	0.57	0.67	0.81	0.38	0.67	0.48	0.59
7 (1a)	0.81	0.87	0.89	0.76	0.85	0.81	0.81
8 (1a)	0.86	0.85	0.89	0.81	0.87	0.83	0.84
9 (1a)	0.74	0.90	0.91	0.73	0.82	0.81	0.81
10 (1a)	0.77	0.85	0.81	0.81	0.79	0.83	0.81
combined (1b)	0.77	0.84	0.81	0.84	0.82	0.84	0.89

highest in AUC, and also in Precision and Recall, whereas persona 6 has poor performance compared to the rest. These results indicate that for high scoring personas, we successfully captured the active CDT campaigns, but for the personas with lower scores, there may not be active campaigns for the period of the experiments.

In order to retrieve the variables that affect the discovery and measurement of CDT, we applied the feature importance method on the dataset of each persona, and selected the top-10 highest scoring features. For the majority of the personas (7 out of 10) the most important features were the number of ads (distinct or not) and the number of keywords in desktop. In some cases, there were also landing pages that had high scoring (i.e., specific ad-campaigns), but this was not consistent across all personas.

Combined Personas: Setup 1b. Here, we use all the datasets collected individually, for each persona in the previous experiment (Setup 1a), and combine them into one unified dataset. This setup emulates the real scenario of a user exhibiting multiple and diverse web interests, that give extra information to the ad-ecosystem about their browsing behavior. Of course, there is an increase in the possible feature space to accommodate all the domains and keywords from all personas. In fact, the dataset contains 2021 features as it stores the vectors of landing pages and keywords, for all the different types of personas. In total, there were 890 distinct ad-domains described by keywords in 76 distinct categories.

In this dataset, we apply feature selection with the Extra-Trees classifier to select the most relevant features and create a more accurate predictive model. This method reduced the feature space to 984 useful features out of 2201. Next, we use the three classification algorithms and a range of hyper-parameters for each one. Also, we apply a 10-fold nested cross-validation method for selecting the best model (in terms of scoring performance) that can give us an accurate, non overly-optimistic estimation [13]. Again, the




Figure 5: Top-30 features ranked by importance using Gini index, in the machine learning model.

best selected model was Random Forest, with 200 estimators (trees) and 200 depth of each tree, with AUC=0.89 (bottom row in Table 2). The model’s performance is high in all the mentioned scores, which indicates that the more diverse data the advertisers collect, the easier it is to identify the different user’s devices. This result is in line with Zimmeck et al. [56], who attempted a threshold-based approach for probabilistic CDT detection on real users’ data, lending credence to our proposed platform’s performance.

We also measure the feature importance for the top-30 features (shown in Figure 5). One third of the top features are related to crawl specific metadata, whereas about half of the top features are keyword-related. Interestingly, features such as the day and time of the experiment, as well as the number of received ads, are important for the algorithm to make the classification of the devices. Indeed, time-related features provide hints on when the ad-ecosystem receives the browsing signal and attempts the CDT, and thus, which days and hours of day the CDT is stronger. These results give support to our initial decision to experiment in a continuous fashion with regular sessions injecting browsing signal, while at the same time measuring the output signal via delivered ads.

5.3 Does long-time browsing improve CDT?

Independent Personas: Setup 2a. In this set of experiments, we allow the devices to train for a longer period of time, to emulate the scenario where a user is focused on a particular interest, and produces heavy browsing behavior around a specific category. This long-lived browsing injects a significantly higher input signal to the ad-ecosystem than the previous setup, which should make it easier to perform CDT. In order to increase the setup’s complexity, and make it more difficult to track the user, we allow all devices (i.e., 1 mobile, 2 desktops) to train in the same way under the same persona. In effect, this setup also tests a basic countermeasure from the user’s point of view, who tries to blur her browsing by injecting traffic of the same persona from all devices to the ad-ecosystem.

In this setup, while all devices are trained with the same behavioral profile, we examine if the statistical tests and ML modeler can still detect and distinguish the CDT. This experiment contains three different phases during each run. The mobile phase, where the mobile performs training crawls for $t_{train}=480$ mins, and a testing crawl for $t_{test}=30$ mins. In parallel with the mobile training, the two desktops perform test crawls for $t_{test}=30$ mins. After mobile training and testing, both desktops start continuous training and testing crawls alternately for 8 hours ($t_{train}=t_{test}=30$ min).

Due to the long time needed for executing this experiment, we focus on two personas constructed in the following way. We use the methodology for persona creation as described in § 4.1, and focus on active ad-campaigns, resulting to two personas in the interest of “Online Shopping-Accessories”, and “Online Shopping-Health and Fitness” (loosely matching the personas 1 and 4 from Table 3). Then, we performed 4 runs of 16 hours duration each, for each persona. In this setup, since all devices are uniformly trained, we do not include the keyword vector of the persona pages into the datasets, to not introduce any bias from repetitive features.

The statistical analysis for this experiment reveals potential CDT, since we accept the null hypothesis for the distribution of ads delivered in the paired desktop and mobile (lowest p-value=0.052), and reject it in the baseline desktop and mobile (highest p-value=0.006). This consistency is interesting, since for this setup all three devices are uniformly trained with the same persona, and thus all of them collect similar ads due to retargeting. However, there is no similarity between the distributions of ads in the devices that do not share the same IP address.

To clarify this finding, we applied the ML algorithms as in the previous experiment. The algorithms again detect CDT between the mobile and the paired desktop, even though all devices were exposed to similar training with the same persona. In fact, Logistic Regression performed the best across both personas, with $AUC \geq 0.81$, and F_1 -score ≥ 0.80 for both classes.⁵ When computing the importance of features, the desktop number of ads and keywords and the desktop time slot are in the top-10 features. Based on these observations, we believe that the longer training time allowed the ad-ecosystem to establish an accurate user profile, and retarget ads on the paired desktop, based on the mobile’s activity.

Combined Personas: Setup 2b. Similarly to § 5.2 we combine all data collected from the Setup 2a into a unified dataset. Under this scenario, in which we mix data from both personas, the classifier again performs well, with AUC=0.89. Important features in this case are the number of ads and keywords delivered to the desktops, the time of the experiment, and number of keywords for the desktop.

Boosted Browsing with CDT trackers and Independent Personas: Setup 2c. In the next set of experiments, we

⁵Detailed evaluation results of § 5.3 presented in Appendix, Table 4.

investigate the role of CDT trackers in the discovery and measurement of CDT. In particular, we attempt to boost the CDT signal, by visiting webpages with higher portion of CDT trackers. Therefore, the experimental setup and the preprocessing method remain the same as in the previous Setup 2a, but we select webpages to be visited that have active ad-campaigns and their landing pages embed the most-known CDT trackers (as we also show in the next section): Criteo, Tapad, Demdex, Drawbridge. We also change the set of our control pages, so that each one contains at least a CDT tracker. News sites have many 3rd-parties compared to other types of sites [22]. Thus, for this boosted browsing experiment, we choose the set of control pages to contain 3 weather pages and 2 news websites,⁶ while verifying they do not serve contextual ads.

Performing the same analysis as earlier, we find that mobile and paired desktop have ads coming from the same distribution (lowest p-value=0.10), and that there is no similarity between the ads delivered in the mobile and baseline desktop (highest p-value=0.007). For a clearer investigation of the importance of the CDT trackers, we also evaluate the findings with the ML models. For persona 1, Logistic Regression and Random Forest models perform near optimally, with high precision of Class 1, high recall for class 0, average F_1 -Score=0.93 for both classes, and AUC=0.93. For persona 4, the scores are even higher, outperforming the other setups, as all metrics for Logistic Regression scored higher than 0.98. Overall, these results indicate that we successfully biased the trackers to identify the emulated user in both devices, and to provide enough output signal (ads delivered) for the statistical algorithms to detect the CDT performed.

Boosted Browsing with CDT trackers and Combined Personas: Setup 2d. We follow a similar approach with before, and combine all data collected from the Setup 2c, into a unified dataset for Setup 2d. Under this scenario, the classifier (Logistic Regression) again performs very well, with AUC=0.93. Important features in this case are the number of ads delivered to the desktops, the time of the experiment in each desktop and the number of keywords. Interestingly, and perhaps unexpectedly, the existence of Criteo tracker in a landing page, is a feature appearing in the top-10 features.

5.4 Does incognito browsing help evade CDT?

Independent Personas: Setup 3a. In this final experimental setup, we investigate if it is possible for the user to apply some basic countermeasures to avoid, or at least reduce the possibility of CDT, by removing her browsing state in every new session. For this, we perform experiments where the traditional tracking mechanisms (e.g., cookies, cache, browsing history, etc.) are disabled or removed, emulating incognito browsing. We select the first five personas from Table 3,

⁶accuweather.com, wunderground.com, weather.com, usatoday.com, huffingtonpost.com

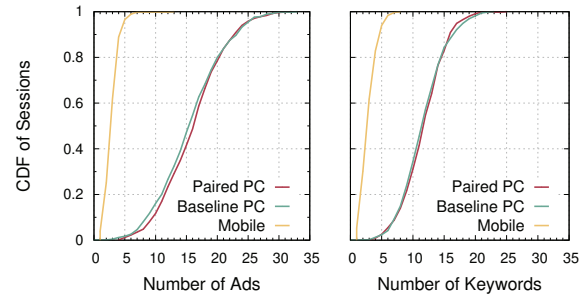


Figure 6: CDF of collected ads (left) and corresponding keywords of the ads (right) per crawling session for all devices.

which had the most active ad-campaigns and appeared to be promising due to the “online shopping” interest. Every desktop executed browsing in a stateless mode, while the mobile in a stateful mode. For each persona, we collected data for two runs, following the timeline of phases as in Setup 1a.

The distributions between mobile vs. paired desktop, as well as mobile vs. baseline desktop, were found to be different (highest p-value=0.034). Also, none of the ML classifiers performed higher than 0.7 (in all metrics), and thus we could not clearly extract any significant result. Specifically, the highest AUC score for personas 1 and 2 was 0.70 with the use of the Random Forest classifier, and for personas 3 and 4 was 0.73 using the Logistic Regression classifier. The worst scoring, independent of algorithm, was recorded for persona 5, with AUC=0.57, and Precision/Recall scores under 0.50.

Combined Personas: Setup 3b. When the data from all five personas are combined, the classifier performing best was Logistic Regression, with AUC=0.79. Overall, these results point to the semi-effectiveness of the incognito browsing to limit CDT. That is, by removing the browsing state of a user on a given device, the signal provided to the CDT entities is reduced, but not fully removed. In fact, when the data from various personas are combined, the CDT is still somewhat effective, since the paired devices have the same IP address.

6 Platform Validation

In this section we validate the representativeness of the data collected from the previous experiments, by examining: (i) the type and frequency of ads delivered in each device, and (ii) the type and number of trackers that our personas were exposed to. We compare the distributions of these quantities with past works and data on real users, to quantify if our synthetic personas successfully emulate real users’ traffic, and if our measurements of the CDT ad-ecosystem are realistic.

We first measure the frequency of ads delivered to our devices in the experiment § 5.2, since it follows a well-crafted timeline that is suitable for this kind of measurement. The ads delivered in the three devices during these sessions are shown in Figure 6 (left). For most sessions (~90%), the mobile device was exposed to fewer than five ads, since the




Figure 7: Top-20 trackers (grouped based on organization) and their coverage in persona pages. For example, all the Google-owned domains, such as Doubleclick, Googleapis, Google-Analytics, are grouped under the “Google” label.

mobile version of websites typically delivers a smaller number of ads, designed for smaller screens and devices. On the contrary, the desktop devices had a higher exposure to ads compared to the mobile device. Also, the two desktops receive a similar number of ads (on average 2 to 4 ads on every visit to the control pages). Similar observations can be made for the keywords categories of ads (Figure 6 (right)). The ad-industry has reported that ~ 300 ads everyday, on average, are being displayed to desktop users [30, 11, 31, 14], while they also recommend the delivery of 5 ads per mobile domain [52], which proportionally match the number of ads we have collected in our mobile and desktop sessions.

We also validate the representativeness of the data collected from the experiments § 5.2 and § 5.3, by examining the trackers appearing in the webpages visited by the personas. We use Disconnect List [18] to detect them and measure their frequency of appearance (i.e., Figure 7). From the trackers detected in the set of persona pages, and using the list provided by [56], 37% was found to be CDT related, including both deterministic and probabilistic. In fact, the top CDT trackers found in our data, which may perform both types of CDT, include Google-owned domains, Facebook, Criteo, Zopim, Bing, Advertising.com(AOL), and are in-line with the top CDT trackers found in [56, 10] (66% overlap of top-20 with [56] and 55% overlap with [10]). In addition, 17% of these trackers are mainly focused on probabilistic CDT, including Criteo, BlueKai, AdRoll, Cardlytics, Drawbridge, Tapad, and each individual tracker is found at least in 2% of the persona pages, again in-line with the results in [56].

7 Discussion & Conclusion

Through extensive experiments with the proposed framework Talon, we were able to trigger CDT trackers into pairing of the emulated users’ devices. This allowed us to statistically verify that CDT is indeed happening, and measure its effectiveness on different user interests and browsing behaviors, independently and in combination. In fact, CDT was

prominent when user devices were trained to browse pages of similar interests, reinforcing the behavioral signal sent to CDT entities, and specifically when browsing activity is related with online shopping, since those types of users seem to be more targeted by advertisers. The CDT effect was further amplified when the visited persona and control pages had embedded CDT trackers, pushing the accuracy of detection up to 99%. We also found that browsing in a stateless mode showed a reduced, but not completely removed CDT effect, as incognito browsing obfuscates somewhat the signal sent to the ad-ecosystem, but not the network access information. Indeed, our data collection was performed across relatively short time periods, in comparison to the wealth of browsing data that advertising networks have at their disposal. In fact, we anticipate that CDT companies collect data about users and devices for months or years, and even buy data from data brokers, to have the capacity of targeting users with even higher rates. To that end, we believe that high accuracies self-reported by CDT companies (e.g., Lotame: $>90\%$ [37], Drawbridge: 97.3% [19]), are possible.

Impact on user privacy: Undoubtedly, CDT infringes on users’ online privacy and minimizes their anonymity. But the actual extent of this tracking paradigm and its consequences to users, the community, and even to the ad-ecosystem itself, are still unknown. In fact, since CDT is heavily depended on user’s browsing activity, and the ad-ecosystem employs such collected data for targeting purposes, one major line of future work is the study of targeting sensitive user categories (e.g., gender, sexual orientation, race, etc.) via CDT. This is especially relevant nowadays with the enforcement of recent EU privacy regulations such as GDPR [24] and ePrivacy [23]. This is where Talon comes in play, as it provides a concrete, scalable and extensible methodology for experimenting with different CDT scenarios, auditing its mechanics and measuring its impact. In fact, the modular design of our methodology allows to study CDT in depth, and propose new extensions to study the CDT ecosystem: new plugins, personas and ML techniques. To that end, our design constitutes Talon into an enhanced transparency tool that reveals potentially illegal biases or discrimination from the ad-ecosystem.

Acknowledgments

The research leading to these results has received funding from the European Union’s Horizon 2020 Research and Innovation Programme under grand agreement No 786669 (project CONCORDIA), the Marie Skłodowska-Curie grant agreement No 690972 (project PROTASIS), and the Defense Advanced Research Projects Agency (DARPA) ASED Program and AFRL under contract FA8650-18-C-7880. The paper reflects only the authors’ views and the Agency and the Commission are not responsible for any use that may be made of the information it contains.

References

- [1] Measuring Cross-Device: The Methodology. <https://www.tapad.com/resources/cross-device/measuring-cross-device-the-methodology>, 2018.
- [2] Pew Research Center - Mobile Fact Sheet. <http://www.pewinternet.org/fact-sheet/mobile/>, 2018.
- [3] ADBRAIN. Demystifying cross-device. essential reading for product management, business development and business technology leaders. <https://www.iabuk.com/sites/default/files/white-paper-docs/Adbrain-Demystifying-Cross-Device.pdf>, 2016.
- [4] ADELPHIC. How cross-device identity matching works. <https://adelphic.com/how-cross-device-identity-matching-works-part-1/>, 2016.
- [5] AGUIRRE, E., MAHR, D., GREWAL, D., DE RUYTER, K., AND WETZELS, M. Unraveling the personalization paradox: The effect of information collection and trust-building strategies on online advertisement effectiveness. *Journal of Retailing* 91, 1 (2015), 34–49.
- [6] ARP, D., QUIRING, E., WRESSNEGGER, C., AND RIECK, K. Privacy threats through ultrasonic side channels on mobile devices. In *IEEE European Symposium on Security and Privacy (EuroS&P)* (2017), pp. 35–47.
- [7] BASHIR, M. A., ARSHAD, S., ROBERTSON, W., AND WILSON, C. Tracing information flows between ad exchanges using retargeted ads. In *25th USENIX Security Symposium* (2016), pp. 481–496.
- [8] BASHIR, M. A., FAROOQ, U., SHAHID, M., ZAFFAR, M. F., AND WILSON, C. Quantity vs. quality: Evaluating user interest profiles using ad preference managers. In *Proceedings of the Annual Network and Distributed System Security Symposium (NDSS), San Diego, CA* (2019).
- [9] BLEIER, A., AND EISENBEISS, M. Personalized online advertising effectiveness: The interplay of what, when, and where. *Marketing Science* 34, 5 (2015), 669–688.
- [10] BROOKMAN, J., ROUGE, P., ALVA, A., AND YEUNG, C. Cross-device tracking: Measurement and disclosures. *Proceedings on Privacy Enhancing Technologies*, 2 (2017), 133–148.
- [11] BRYCE SANDERS. Do we really see 4,000 ads a day? <https://www.bizjournals.com/bizjournals/how-to/marketing/2017/09/do-we-really-see-4-000-ads-a-day.html>, 2017.
- [12] CARRASCOSA, J. M., MIKIANS, J., CUEVAS, R., ERRAMILLI, V., AND LAOUTARIS, N. I always feel like somebody’s watching me: measuring online behavioural advertising. In *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies (CONEXT)* (2015).
- [13] CAWLEY, G. C., AND TALBOT, N. L. On over-fitting in model selection and subsequent selection bias in performance evaluation. *Journal of Machine Learning Research* 11 (2010).
- [14] CHRISTOPHER ELLIOTT. Yes, there are too many ads online. yes, you can stop them. heres how. https://www.huffingtonpost.com/entry/yes-there-are-too-many-ads-online-yes-you-can-stop_us_589b888de4b02bbb1816c297, 2017.
- [15] CRITEO. The State of Cross-Device Commerce. <https://www.criteo.com/wp-content/uploads/2017/07/Report-criteo-state-of-cross-device-commerce-2016-h2-SEA.pdf>, 2016.
- [16] CRITEO. The 5 top attribution methodologies for cross-channel roi. <https://www.criteo.com/insights/top-attribution-methodologies-for-cross-channel-roi/>, 2018.
- [17] DAS, A., BORISOV, N., AND CAESAR, M. Tracking mobile web users through motion sensors: Attacks and defenses. In *Proceedings of the Annual Network and Distributed System Security Symposium (NDSS), San Diego, CA* (2016).
- [18] DISCONNECT. Disconnect lets you visualize and block the invisible websites that track your browsing history. <https://disconnect.me/>, 2019.
- [19] DRAWBRIDGE. Cross-Device Consumer Graph. <https://go.drawbridge.com/rs/454-0RY-155/images/Drawbridge-Cross-Device-Consumer-Graph.pdf>, 2015.
- [20] DRAWBRIDGE. Drawbridge Cross-Device Connected Consumer Graph Is 97.3% Accurate. <https://go.drawbridge.com/rs/454-0RY-155/images/Drawbridge-Cross-Device-Consumer-Graph.pdf>, 2015.
- [21] EASYLIST. Easylist is the primary filter list that removes most adverts from international webpages, including unwanted frames, images and objects. <https://easylist.to/>, 2018.

- [22] ENGLEHARDT, S., AND NARAYANAN, A. Online tracking: A 1-million-site measurement and analysis. In *Proceedings of the ACM SIGSAC conference on computer and communications security (CCS)* (2016), pp. 1388–1401.
- [23] EUROPEAN PARLIAMENT, COUNCIL OF THE EUROPEAN UNION. Directive 2002/58/EC of the European Parliament and of the Council of 12 July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector (Directive on privacy and electronic communications), 2002.
- [24] Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). *Official Journal of the European Union L119* (2016), 1–88.
- [25] FARAHAT, A., AND BAILEY, M. C. How effective is targeted advertising? In *Proceedings of the 21st ACM International Conference on World Wide Web* (2012), pp. 111–120.
- [26] GIRONDA, J. T., AND KORGAONKAR, P. K. ispy? tailored versus invasive ads and consumers perceptions of personalized advertising. *Electronic Commerce Research and Applications* 29 (2018), 64–77.
- [27] GOOGLE. Google Product Taxonomy. <https://www.google.com/basepages/producttype/taxonomy.en-US.txt>, 2015.
- [28] GOOGLE. Run apps on the Android Emulator. <https://developer.android.com/studio/run/emulator/>, 2018.
- [29] IORDANOU, C., KOURTELLIS, N., CARRASCOSA, J. M., SORIENTE, C., CUEVAS, R., AND LAOUTARIS, N. Beyond content analysis: Detecting targeted ads via distributed counting. *arXiv preprint arXiv:1907.01862* (2019).
- [30] JON SIMPSON. Finding brand success in the digital world. <https://www.forbes.com/sites/forbesagencycouncil/2017/08/25/finding-brand-success-in-the-digital-world/#734eaba626e2>, 2018.
- [31] JUSTIN MALLINSON. How many ads do we really see each day? <http://www.tcsmedia.co.uk/many-ads-really-see-day/>, 2018.
- [32] KAFKA, P., AND MOLLA, R. Recode - 2017 was the year digital ad spending finally beat TV. <https://www.recode.net/2017/12/4/16733460/2017-digital-ad-spend-advertising-beat-tv>, 2017.
- [33] KOROLOVA, A., AND SHARMA, V. Cross-app tracking via nearby bluetooth low energy devices. In *Proceedings of the 8th ACM Conference on Data and Application Security and Privacy (CODASPY)* (2018), pp. 43–52.
- [34] LÉCUYER, M., DUCOFFE, G., LAN, F., PAPANCEA, A., PETSIOS, T., SPAHN, R., CHAINTREAU, A., AND GEAMBASU, R. Xray: Enhancing the webs transparency with differential correlation. In *23rd USENIX Security Symposium* (2014), pp. 49–64.
- [35] LEWIS, R. A., RAO, J. M., AND REILEY, D. H. Here, there, and everywhere: correlated online behaviors can lead to overestimates of the effects of advertising. In *Proceedings of the 20th ACM International Conference on World Wide Web* (2011), pp. 157–166.
- [36] LIU, B., SHETH, A., WEINSBERG, U., CHANDRASHEKAR, J., AND GOVINDAN, R. Adreveal: improving transparency into online targeted advertising. In *Proceedings of the 12th ACM Workshop on Hot Topics in Networks* (2013), p. 12.
- [37] LOTAME. Cross-Device ID Graph Accuracy: Methodology. <https://www.lotame.com/cross-device-id-graph-accuracy-methodology/>, 2016.
- [38] LOTAME. Cross-device bridging the gap between screens. <https://www.lotame.com/products/cross-device/>, 2018.
- [39] MATHUR, A., VITAK, J., NARAYANAN, A., AND CHETTY, M. Characterizing the use of browser-based blocking extensions to prevent online tracking. In *Fourteenth Symposium on Usable Privacy and Security (SOUPS)* (Baltimore, MD, 2018), USENIX Association, pp. 103–116.
- [40] MAVROUDIS, V., HAO, S., FRATANONIO, Y., MAGGI, F., KRUEGEL, C., AND VIGNA, G. On the privacy and security of the ultrasound ecosystem. *Proceedings on Privacy Enhancing Technologies* (2017).
- [41] MCAFEE. Customer URL Ticketing System. <https://www.trustedsource.org/>, 2018.
- [42] MCNAIR, C. Global Ad Spending Update. <https://www.emarketer.com/content/global-ad-spending-update>, 2018.
- [43] MELICHER, W., SHARIF, M., TAN, J., BAUER, L., CHRISTODORESCU, M., AND LEON, P. G. (do not) track me sometimes: Users contextual preferences for web tracking. *Proceedings on Privacy Enhancing Technologies*, 2 (2016).

- [44] PACHILAKIS, M., PAPADOPOULOS, P., MARKATOS, E. P., AND KOURTELLIS, N. No more chasing waterfalls: A measurement study of the header bidding ad-ecosystem. In *19th ACM Internet Measurement Conference* (2019).
- [45] PAN, X., CAO, Y., AND CHEN, Y. I do not know what you visited last summer - protecting users from third-party web tracking with trackingfree browser. In *Proceedings of the Annual Network and Distributed System Security Symposium (NDSS), San Diego, CA* (2015).
- [46] PAPADOPOULOS, E. P., DIAMANTARIS, M., PAPADOPOULOS, P., PETSAS, T., IOANNIDIS, S., AND MARKATOS, E. P. The long-standing privacy debate: Mobile websites vs mobile apps. In *Proceedings of the 26th ACM International Conference on World Wide Web* (2017), pp. 153–162.
- [47] PAPADOPOULOS, P., KOURTELLIS, N., AND MARKATOS, E. Cookie synchronization: Everything you always wanted to know but were afraid to ask. In *The World Wide Web Conference* (2019), ACM, pp. 1432–1442.
- [48] PAPADOPOULOS, P., KOURTELLIS, N., RODRIGUEZ, P. R., AND LAOUTARIS, N. If you are not paying for it, you are the product: How much do advertisers pay to reach you? In *Proceedings of the ACM Internet Measurement Conference* (2017), pp. 142–156.
- [49] PATRICK HOLMES. Mobile and Desktop Advertising Strategies Based on User Intent. <https://instapage.com/blog/adwords-search-device-user-intent>, 2018.
- [50] PROJECT, J. F. Automation for Apps. <http://appium.io/>, 2018.
- [51] RAMIREZ, E., OHLHAUSEN, M., AND MCSWEENEY, T. Cross-device tracking: An FTC staff report. Tech. rep., 2017.
- [52] RENE HERMENAU. Adsense max allowed number of ads - 2018 Rules. <https://wpquads.com/google-adsense-allowed-number-ads/>, 2018.
- [53] TAPAD. Tapad device graph - creating a unified view of the consumer. <https://www.tapad.com/device-graph/>, 2018.
- [54] TAPAD. The expert’s guide to cross-device conversion & attribution. <https://www.tapad.com/uses/the-experts-guide-to-cross-device-conversion-attribution>, 2018.
- [55] TUCKER, C. E. Social networks, personalized advertising, and privacy controls. *Journal of Marketing Research* 51, 5 (2014), 546–562.
- [56] ZIMMECK, S., LI, J. S., KIM, H., BELLOVIN, S. M., AND JEBARA, T. A privacy analysis of cross-device tracking. In *26th USENIX Security Symposium* (2017), pp. 1391–1408.

A Appendix

Table 3: Behavioral personas used in our experiments.

Persona	Category - Description
1	Online Shopping - Accessories, Jewelry.
2	Online Shopping - Fashion, Beauty.
3	Online Shopping - Sports and Accessories.
4	Online Shopping - Health and Fitness.
5	Online Shopping - Pet Supplies.
6	Air Travel.
7	Online Courses and Language Resources.
8	Online Business, Marketing , Merchandising.
9	Browser Games - Online Games.
10	Hotels and Vacations.

Table 4: Performance evaluation for Logistic Regression in total components of Setup 2. Left value in each column is the score for Class 0 (C0=not paired desktop); right value for Class 1 (C1=paired desktop).

Persona (setup)	Precision		Recall		F_1 -Score		AUC
	C0	C1	C0	C1	C0	C1	
1 (2a)	0.90	0.79	0.82	0.88	0.86	0.83	0.85
4 (2a)	0.83	0.79	0.81	0.81	0.82	0.80	0.81
combined(2b)	0.87	0.92	0.92	0.87	0.89	0.90	0.89
1 (2c)	0.87	1.0	1.0	0.88	0.93	0.93	0.93
4 (2c)	1.0	0.98	0.98	1.0	0.99	0.99	0.99
combined(2d)	1.0	0.86	0.88	1.0	0.93	0.93	0.93

Table 5: Description of features used by datasets. The type of desktop crawl values are in range $\{0,1\}$, where 0 represents the before/test sessions, while 1 the after/train sessions. The time of crawl is divided in 30 minutes timeslots and is encoded in range $\{0,48\}$. The day of crawl is encoded in range $\{1,7\}$. V represents the (enumerated) vectors of values in the sets of: landing pages, training pages, ads and ad categories.

Feature Label	Description
Crawl.Type	The type of desktop crawl.
Run_ID	The indexed number of run $\{1,4\}$.
Session_ID	The index of session $\{1,15\}$.
Persona_Keywords	V: keyword categories of training pages.
Mobile_Timeslot	Time of crawl (Mobile).
Desktop_Timeslot	Time of crawl (Desktop).
Desktop_Day	The day of crawl (Desktop).
Mobile_Number_of_Ads	# ad domains (Mobile).
Desktop_Number_of_Ads	# ad domains collected (Desktop).
Mobile_Unique_Number_of_Ads	# distinct ad domains (Mobile).
Desktop_Unique_Number_of_Ads	# ad domains (Desktop).
Mobile_Number_of_Keywords	# ad categories (Mobile).
Desktop_Number_of_Keywords	# ad categories (Desktop).
Mobile_Unique_Number_of_Keywords	# distinct ad categories (Mobile).
Desktop_Unique_Number_of_Keywords	# distinct ad categories (Desktop).
Mobile_Keywords	V: keyword categories of landing pages (Mobile).
Desktop_Keywords	V: keyword categories of landing pages (Desktop).
Mobile_Landing_Pages	V: landing pages of delivered ads (Mobile).
Desktop_Landing_Pages	V: landing pages of delivered ads (Desktop.)