

Pay-per-Tracking: A Collaborative Masking Model for Web Browsing

Javier Parra-Arnau

Department of Computer Engineering and Mathematics, Universitat Rovira i Virgili, Av. Països Catalans 26, E-43007 Tarragona, Catalonia, e-mail javier.parra@urv.cat.

Abstract

Web tracking is the key enabling technology of modern online advertising and, at the same time, the source of serious privacy concerns. In recent years, we have witnessed the emergence of a variety of technologies whose main goal is to address these concerns. However, ad blockers and anti-trackers eliminate all forms of tracking and advertising and therefore fail to reconcile user privacy and the current Internet business model.

In this paper, we propose a new tracking paradigm that aims at returning control to users over tracking and advertising, and allowing them to participate in the monetization of their browsing data. The proposed paradigm breaks with the current barter model of exchanging privacy for services and, at the same time, it may preserve the Internet economic model where content is paid from advertisement revenue. We design a system architecture that implements this model in practice, and optimizes the exchange of privacy for money. Our ultimate aim is to strike a better balance between user privacy and the economic model that sustains the Web, and thus overcome the deadlock caused by the ad blocking wars. Experimental results with real browsing data demonstrate the suitability and feasibility of our approach.

Keywords: user privacy, Web tracking, user profiling, collaborative masking, privacy-utility trade-off

1. Introduction

With a revenue of \$27.5 billion in the first half of 2015, the online advertising industry is one of the most competitive industries in the world today, and an example of the transformation driven by the ever-growing sophistication of personalization technologies [11]. In the past, ads were served directly by the Web site's owner following a one-size-fits-all approach. But due to the gradual introduction of intermediary companies with extensive capabilities to track users, Internet advertising has become increasingly customized and pervasive.

Although ads can be targeted to several aspects of a user (e.g., location and page content), those relying on their browsing interests are by far the most effective form of advertising, ensuring conversion rates¹ that double those of geographical and contextual ads [31]. However, *behavioral targeting* —as it is also known— poses the greatest threat to privacy, since it is the result of tracking and profiling users' browsing habits throughout the Internet, often without their knowledge [68] and consent².

The intrusiveness of these tracking and profiling practices, as well as the increasing invasiveness of digital advertising, have raised serious concerns regarding user privacy and Web usability in the last years. According to recent studies, two out of three Internet users are worried about the fact that their online behavior be scrutinized by an industry that operates in the shadows with virtually no oversight. Numerous surveys in this line reflect the growing level of ubiquity and abuse of advertising, which is perceived by users as a significant degradation of their browsing experience [18, 62].

As a reaction to these privacy and usability concerns, we have recently witnessed the emergence of a wide variety of tools whose main goal is to block ads and avoid tracking. Examples of these tools include Adblock

¹In online marketing terminology, conversion usually means the act of converting Web site visitors into paying customers.

²Consistently with the recommendations of the US Federal Trade Commission, the advertising industry has started to offer an opt-out scheme for behavioral advertising [5].

Plus [1] and Ghostery [9], the former being the most downloaded browser plug-in for Google Chrome and Safari worldwide. However, ad blockers and anti-trackers constitute a radical approach in this bid to regain control over online advertising and tracking. While they may make Web sites cleaner and faster, users can *only* choose between blocking or allowing tracking, and hence, between blocking or allowing ads³. Nowadays, with more than 200 million people worldwide regularly using these technologies, the current Internet business model—in which users get content free in return for receiving ads and giving up some personal data—is at serious risk. Just in 2015, ad blocking was estimated to cost publishers nearly \$22 billion [21].

The progressive transformation of this model towards a paradigm where users access content for nothing in return is leading the online advertising industry to adopt new strategies. Big advertising firms have started to pay ad-blockers’ companies for their ads to be unblocked [33], and some publishers have engaged in an arms race and started to deny access to users of these tools [65]. All this has spurred a heated debate about the ethics of ad blockers, confronting their users with the advertising industry in what has been called “the ad blocking wars” [66].

At the same time, this situation is fostering new business opportunities with respect to the monetization of personal data, shifting the balance of power between users and the companies that harvest said information. A number of start-ups led by DataCoup [6], DataWallet [7] and CitizenMe [2] has recently embarked on this monetization project and, over the course of this year, will allow users to sell their social networks data directly to businesses and advertisers [35, 46]. In essence, these firms will operate as data brokers and enable users to negotiate deals with parties interested in their accounts on Facebook, Twitter, Google+, LinkedIn and related social networks. In this same monetization spirit, Facebook is now considering paying users to post on its site [67].

1.1. Contribution and Plan of this Paper

In this paper, we propose a new tracking model whereby users can regain full control over Web tracking and advertising. Following the recent tendency of some companies to pay for online user information, we propose that users exert this control through the monetization of their browsing profiles.

In our scheme, users’ tracking preferences are enforced in a flexible manner, unlike the simple, binary choices provided by ad blockers and anti-trackers. The monetization of tracking, on the other hand, is optimized so that user privacy is maximized for an economic reward or compensation an ad company may offer.

The proposed tracking scheme may not be regarded as a radical shift in this effort to return control to users, in the sense that, if an ad company is not willing to pay for tracking, then it is driven out of the business. In fact, our solution allows companies to continue tracking and serving ads on a fraction of the pages they are present, even though they do not bid for tracking. However, if users are not rewarded, the browsing profiles built from the tracked pages do not capture real interests and, hence, the ads personalized to these interests are no longer effective.

Our proposal therefore breaks with the current barter model of exchanging privacy for services and, at the same time, it may preserve the Internet business model where content is paid from advertisement revenue. With it, our ultimate aim is to strike a better balance among user privacy and this business paradigm, and thus overcome the deadlock caused by the ad blocking wars.

Next, we summarize the major contributions of this work:

- We propose a tracking model that enables users to control and monetize the disclosure of private browsing data to ad platforms and trackers. Unlike current privacy technologies such as Adblock Plus and Ghostery, which offer the possibility of either fully revealing or completely obfuscating user-profile information, our solution enables users to disclose a masked version of their browsing profiles, still useful for ad personalization. Building upon the principle of co-utility [36], we assume that users will collaborate among themselves so as to make an informed decision about the level of masking to be applied. Accordingly, we formulate the problem of perturbing a user profile as a multiobjective optimization problem that takes into account privacy and money.

³The vast majority of ads today are served by third-party trackers [86, 20].

- We design a system architecture implementing our pay-per-tracking model, and specify the fundamental operational structure of its main component, a Web-browser plug-in. Installed on a user’s machine, the plug-in perturbs their browsing data in real-time and in an automated fashion, by selectively blocking and allowing tracking on the pages they visit. We propose a masking algorithm to this end, which achieves serviceable points of operation within the privacy-money trade-off, and which causes the least possible impact on the Web economy.
- We conduct a thorough experimental analysis to show the technical feasibility of our tracking scheme, and the benefits it would provide to both users and ad companies. Our experimental evaluation relies on the browsing data and ads of 144 users, and investigates numerous elements of Web tracking and advertising, including the enhanced Web presence of ad companies, the effectiveness of profile-based ads, and the impact of our solution on user privacy and the Internet economy.

The remainder of this work is organized as follows. Sec. 2 provides the necessary background in online advertising and reviews the state of the art relevant to this work. Then, Sec. 3 presents our pay-per-tracking model, a collaborative masking mechanism for the disclosure of browsing data to ad companies, and a formulation of the trade-off between privacy and money. Sec. 4 describes the main components of a system architecture implementing our solution. Sec. 5 conducts an experimental evaluation of the proposed tracking model. Conclusions are drawn in Sec. 6. Finally, Appendix Appendix A shows the centroids of the profile clusters used in the experimental section.

2. Background and Related Work

In this section, we first explore the online advertising ecosystem, providing the reader with the necessary depth to understand the technical contributions of this paper. Secondly, we review the state of the art relevant to this work.

2.1. Background in Online Advertising

In this section, we first give an overview of the main actors of the advertising ecosystem. Afterwards, we describe how ads are served on the Web, and then, we provide a standard classification of the targeting objectives commonly available to advertisers. For a detailed, complete explanation on the subject, the reader is referred to [86].

2.1.1. Key Actors

The online advertising industry is composed by a considerable number of entities with very specific and complementary roles, whose ultimate aim is to display ads on Web sites. Publishers, advertisers, ad platforms, ad agencies, data brokers, aggregators and optimizers are some of the parties involved in the ad-delivery process [96]. Despite the enormous complexity⁴ and constant evolution of the advertising ecosystem, the process whereby ads are presented on Web sites is usually characterized or modeled in terms of publishers, advertisers and ad platforms [89, 58, 93, 28, 91]. Next, we provide a description of these three key actors:

- A *publisher* is an entity that owns a Web page (or a Web site) and that, in exchange of some economic compensation, is willing to place ads of other parties in some spaces of its page (or site). An example of publisher is The New York Times’ Web site.
- An *advertiser* is an entity that wants to display ads on one of the spaces offered by a publisher, and is disposed to pay for it. Advertisers typically engage the services of one or several *ad platforms* (described below), which are the ones responsible for displaying their ads on the publishers’ sites. Among other requirements, advertisers indicate the targeting objective/s most suitable for their campaigns, that is,

⁴ The intricacy of the advertising ecosystem is often illustrated in conferences and related venues with the diagram available at [53].

to which users they want their ads to be shown. For example, an advertiser may want the ad platform to serve its ads to an audience interested in politics or to people living in New York City. Advertisers must also specify the amount of money they are willing to pay each time their ads are displayed, and each time users click on them⁵.

- An *advertising platform* or *ad platform* is a group of entities that connects advertisers to publishers, i.e., it receives ads from advertisers and places them on the spaces available at publishers. To this end, ad platforms track and profile users with the aim of targeting ads to their interests, location and other personal data. Ad platforms usually carry out this targeting on their own, in accordance with the campaign requirements and objectives specified by advertisers. Some examples of ad platforms include DoubleClick, Gemini and Bing Ads, owned respectively by Google, Yahoo! and Microsoft.

2.1.2. Ad-Serving Process

Without loss of rigor, throughout this work we shall assume an online advertising model composed mainly of the three entities set forth above. In this simplified albeit comprehensive terms, the ad-delivery process begins with publishers embedding in their sites a link to the ad platform/s they want to work with. The upshot is as follows: when a user retrieves one of those Web sites and loads it, their browser is immediately directed to all the embedded links. Then, through the use of third-party cookies, Web fingerprinting or other tracking technologies, the ad platform is able to track the user's visit to this and any other site partnering with it.

As one might guess, the ability of tracking users across the Web is of paramount importance for ad platforms: it enables them to learn the Web page being visited and hence its content; the user's location through their IP address; and, more importantly, their Web-browsing interests; we note that ad platforms may also acquire tracking data from other entities like data brokers, data management platforms and specialized tracking companies. Afterwards, all these invaluable data about the user, collected by themselves or not, is what allow ad platforms to serve *targeted* ads.

To carry out this task, the vast majority of ad platforms rely on proprietary *targeting algorithms* [86]. The aforementioned user data and the objectives and budgets of all advertisers for displaying their ads are the inputs of these algorithms, which are responsible for selecting which ad will be shown in a particular ad space. Evidently, their primary aim is to maximize ad-platforms' revenues whilst satisfying advertisers' demand.

Finally, the ad-serving process ends up by displaying the selected ad in the user's Web browser, a last step that may entail a content-delivery network. Fig. 1 shows the current architecture of online advertising composed mainly by publishers, ad platforms and advertisers, and illustrates the process whereby third-party ads are displayed to users.

Last but not least, we would like to stress that the advertising model described here —and examined in this paper— corresponds to *indirect-sale advertising*, also called network-based or third-party advertising. This is in contrast to the *direct-sale advertisement* model, where publishers and advertisers negotiate directly, without the mediation of ad platforms. In this latter case, we mostly find popular Web sites selling ad space directly to large advertisers. The ads served this way are essentially untargeted, and are often displayed in Web sites where the products and services advertised are related to their contents. This is mainly because the capability of a publisher to track and profile users is limited just to its own Web site and maybe a few partners. For example, the New York Times' Web site may track users also across the International Herald Tribune, owned by its media group. Such a tracking capability, however, is ridiculous when we compare it with the 2 million sites reachable by Google's ad platforms [30].

2.1.3. User-Targeting Objectives

The ads delivered through indirect-sale advertising allow advertisers to target different aspects of a Web user. In this subsection, we briefly review the most popular *targeting objectives*, which include serving ads

⁵In the terminology of online advertising, these quantities are referred to as the cost-per-impression (CPI) and the cost-per-click (CPC), respectively.

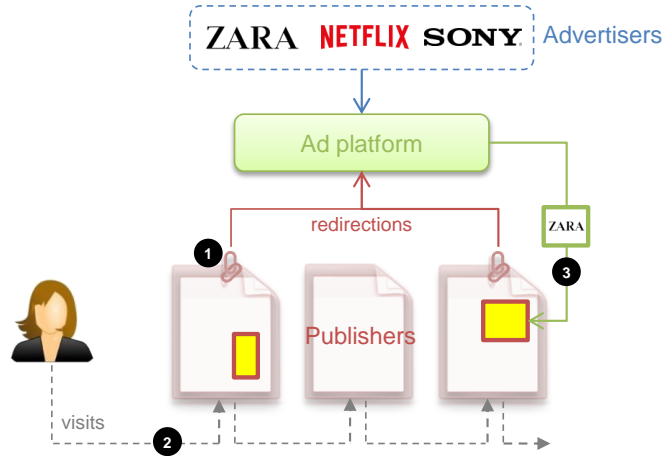


Figure 1: Current online advertising architecture composed by publishers, ad platforms and advertisers. The ad-delivering process requires that publishers include a link to the ad platform they want to partner with (1); for the sake of simplicity, we consider here a single ad platform. When a user visits pages partnering with this ad platform, the browser is instructed to load the URLs provided by the ad platform. Through the use of third-party cookies and other tracking mechanisms, the ad platform is able to track all these visits and build a browsing profile (2). Based on this profile, the user’s location and other parameters, the ad platform uses its targeting algorithm to decide which ad to present on the publisher’s page.

tailored to the Web page they are currently visiting, their geographic location, and their Web-browsing interests.

- *Contextual ads.* Advertisers can reach their audience through contextual and semantic advertising, by directing ads related to the content of the Web site where they are to be displayed. An example of such targeting strategy would be a health-insurance company wishing to show their ads in Web sites whose content is classified as “health & fitness”.
- *Location-based ads.* They are generated based on the user’s location, for example, given by the GPS of their smartphone or tablet, and also according to the IP address of the user’s machine or device. Geographically-targeted ads enable advertisers to launch campaigns targeting users within a certain geographical location.
- *Interest-based or profile-based ads.* Advertisers can also target users based on their Web-browsing interests. These interests are inferred from the *pages tracked* by ad platforms and other tracking companies that may share this information with the former. The sequence of Web sites browsed by a user and effectively tracked by an ad platform or tracker is referred to as the user’s *clickstream*. In current practice, this is the information leveraged by the online advertising industry to construct a user’s interest profile [3, 90, 15, 58, 93, 28, 72, 86].
- *Generic ads.* Advertisers can also specify ad placements or sections of publisher’s Web sites (among those partnering with the ad platform) where their ads will be displayed. Ads served through placement targeting are not necessarily in line with the Web site’s content, but may simply respond to some match between the interests of the visiting users and the products advertised. Because these ads do not rely on any user data, we shall also refer to them as *generic ads*.

An important aspect of the targeting objectives described above is that the former three are not mutually exclusive. In other words, except for placement ads—which are considered to be untargeted—, ads can be simultaneously directed based on content, location and interests.

2.2. Related Work

We proceed by exploring, first, the current software technologies aimed at blocking ads; and secondly, we examine those initiatives intended to address the ad blocking wars.

As mentioned in the introductory section, the vast majority of the content on the Web is sustained by online advertising. In the Web’s predominant economic model, users get content free in return for allowing advertisers to show them ads. However, ads come at the cost of Web-browsing experience, mobile data charges [27], and more importantly, due to the greater effectiveness of profile-based ads⁶, they come at the expense of extensive user tracking and profiling.

To address these usability and privacy concerns, a myriad of ad blockers and anti-trackers have emerged in the last years. In essence, these tools act as firewalls between the browser on the one hand, and the ad platforms and tracking companies on the other. In current practice, there is no difference between ad blockers and anti-trackers. Typically, both technologies monitor all network connections that may be initiated when the browser loads a page, and prevent those which are made with third parties and may entail the delivery of an ad⁷. These network connections are commonly referred to as *third-party network requests* and, as explained in Sec. 2.1, allow their requesters to track users through third-party cookies, Web fingerprinting and other technologies. Because the operation of ad blockers therefore boils down to preventing said tracking, henceforth we shall refer to them simply as anti-trackers.

Most anti-trackers are implemented as open-source browser plug-ins, and block tracking with the help of a data base or *blacklist* of ad platforms and trackers. Basically, these lists include regular expressions and rules to filter out the third-party network requests that are considered to belong to ads or trackers. The maintenance of such blacklists is done manually by the technologies’ developers and in some cases by user communities. Among the most popular anti-tracking technologies, we find Adblock Plus [1], Ghostery [9], Disconnect [8], Lightbeam [12] and Privacy Badger [14], which, from an operational point of view, are equivalent.

The problem with anti-trackers is that they are extremely limited and radical in their approach. Essentially, they either block tracking or allow it throughout the Web⁸. To limit the impact of ad blocking on the advertising industry and the Web, several strategies have been adopted in the last years. Among the most popular ones, a middle-ground approach for ad-blocking has emerged that uses whitelists to allow only “acceptable ads”. The criteria for acceptability typically comprise non-invasiveness, silence and small size [82]. However, because these criteria ultimately depend on the anti-trackers’ developers, this approach does not signify any real advance in the direction of returning users control over tracking and advertising. Indeed, this “acceptable-ads” approach has caused a great controversy in the industry, when it came to the public that the most popular ad blocker was accepting money from some of the whitelisted companies [33].

The Interactive Advertising Bureau (IAB) [10], one of the biggest advertising organizations with over 5 500 ad companies and publishers, has recently released a guide to actively fight ad blocking. The guide is aimed at publishers, and recommends a series of measures that range from restricting access to content, to asking visitors to pay, and to rewarding visitors for interacting with ads. We believe, however, that these recommendations fail to tackle the root of the problem, that is, the reasons why users end up installing anti-trackers. According to a recent survey, two out of three ad-blocker users are not against ads and would accept the trade-off that comes with the “free” content [17]. Nonetheless, this is provided that advertising is a transparent process and they have control over the personal information that is collected [80].

Lastly, we would like to mention Your Online Choices [95] and Do Not Track (DNT) [23], two attempts by the Internet advertising industry and the World Wide Web Consortium to address the aforementioned privacy issues. Although these two self-regulatory initiatives make opt-out easier for users —the former to stop receiving ads tailored to their Web-browsing interests, and the latter to stop being tracked through third-party cookies—, the fact is that users have no control over whether or not their advertising and tracking preferences are honored.

⁶According to a 2010 US report by marketers NAI [31], profile-based advertising “is more than twice as effective at converting users who click on the ads into actual buyers”.

⁷It is worth noting that this indirect form of advertising accounts for the vast majority of ads today [86, 20].

⁸Anti-trackers also allow users to conduct said blocking on a per-page basis, upon users’ request. For example, a user may not feel comfortable with being tracked on a certain Web page and may decide preventing it only in this page.




Figure 2: The perturbation of browsing data, in the form of forgery, suppression, generalization, clustering and other transformation procedure, allows a user to expose a masked version of this private information, to obtain a certain economic compensation, at the expense of some loss in privacy.

3. Pay-per-Tracking

In this section, we investigate a new tracking model that aims at returning control to users over tracking and advertising, and allowing them to participate in the monetization of their browsing data.

The proposed model does not require the complete re-consideration of the prevailing advertising infrastructure, does not alter the current ad-delivery process, nor has the devastating effect of ad blockers and anti-trackers on the Web economy. Rather, our approach may preserve to a certain extent the current Internet business model based on advertising.

Recall from Sec. 2.1.1 that ad platforms track users directly on their own, and/or indirectly through specialized tracking companies. From a technical perspective, although both are regarded as ad companies within the online advertising ecosystem, the difference between an ad platform and a tracker is that the latter does not serve ads to users. Throughout this section, we shall maintain this distinction but, occasionally and for simplicity, we may refer to ad platforms just as trackers. Whether the term “tracker” means “ad platform” or not will be clear from the context.

3.1. A New Online Tracking Model

As explained in the background section, Web tracking is key to the online advertising industry. Among the four classes of ads described in Sec. 2.1.3, profile-based ads are the most effective form of ad targeting, providing conversion rates that double those of location-based and contextual ads [31]. However, behavioral advertising, the targeting of advertisements based on a user’s Web browsing, remains a major source of privacy invasion since it implies tracking and profiling users.

Our proposal targets users who are not in general against advertising and tracking, and understand and appreciate the crucial role of online advertising as the major sustainer of the Internet “free” services. We consider that they may be, in fact, disposed to be tracked and profiled, but as long as their privacy preferences are respected and they have control over the browsing data collected by ad platforms and trackers.

To address the current state of affairs posed by ad blocking, we propose a new tracking model in which these users, instead of blocking all tracking, are willing to share portions of their browsing profiles, or modified versions of it, with such companies.

In our tracking model, we assume users take charge of this *profile masking or perturbation*, given their lack of trust in any external entity—not to mention ad platforms or trackers—to protect said sensitive information. However, an immediate question arises when considering the disclosure of personal data, even completely masked: why a user should take the trouble to generate a perturbed version of their profile (with the potential risks it may entail), if their privacy is fully preserved with the simple use of an ad blocker?

In this work, we address this question by introducing a strong motivation for users to hand this information over, namely, the *monetization of tracking*. Although we shall assume a direct economic reward for exposing their profiles, maybe through micropayments, we would like to stress that other forms of compensation like coupons and discounts might of course be possible.

In our model, the level of perturbation applied to the data and hence the intensity of the data-protection method will depend on this economic compensation. One could argue, though, that users are already recompensed for being tracked: they have access to publishers’ content and services. Nevertheless, if users lose control over the information that is collected on the Web and therefore cannot evaluate the impact that

tracking may have on their privacy, users can hardly determine if those services are worth the relinquished privacy.

The proposed tracking model precisely breaks with the current barter scheme of exchanging privacy for services, which has led us to the present situation of privacy abuse. In a monetary economy, money plays the role of a measure of value of all goods, so that their values can be assessed against each other [29]; this role is absent in a barter economy.

Our proposal provides a more efficient and fairer model where users can monetize their personal data and decide if the offered services are worth the money or not. An example of such monetization might be a user accepting low fares for being tracked on news-related Web sites, and this same user imposing higher prices on sensitive Web-page categories like health and religion.

With this new scheme, we also make the first attempt to incorporate users into the online advertising business. Publishers, advertisers and ad platforms all earn money by tracking users. But in this current model, users, who are the indispensable centerpiece of the business, do not only participate in it but also their privacy is jeopardized.

We would like to stress that the proposed user-centric scheme must not be interpreted as a radical shift in this effort to regain control over tracking and advertising, in the sense that, if an ad platform does not want to compensate a user for tracking, then it is driven out of the business. As we shall explain later in greater detail, ad platforms and trackers can continue tracking and serving ads on a fraction of the pages they are present, even though they do not bid for tracking. However, the browsing profiles built from the tracked pages will not reflect real interests and, consequently, the interest-based ads served this way will be meaningless.

Lastly, we emphasize the suitability of the proposed monetization model with a couple of recent surveys [22, 57] that indicate that more than three out of five Internet users would be willing to exchange their online private data for an economic reward.

3.2. Collaborative Masking

The primary aim of the tracking model presented in Sec. 3.1 is to return control to users over tracking and thus advertising. In this bid to regain control, we shall assume users are reluctant to delegate the protection of their privacy to external entities, and hence opt for taking an active role in safeguarding their browsing data.

With this assumption in mind, we propose a *collaborative, distributed masking* scheme that builds on the idea first developed by Soria-Comas et al [87] for statistical disclosure revelation [48]. Collaborative masking may be regarded as an enhancement of local perturbation, a common approach to data protection when users mistrust any communicating entity, and decide to take their own responsibility for preserving their privacy. The fundamental principle of this local approach is that each user protects their data themselves and *independently* of the other users. Although this eliminates the need for trusted entities or peers, local masking suffers from two major limitations that impede its application to the tracking model at hand.

On the one hand, because users mask their data without seeing the other users' data, they are unable to determine the amount of perturbation that yields a good trade-off between privacy risk and the benefits of accurate data [87]. A natural tendency is for each user to play it safe and overdo the perturbation just in case, which incurs more data inaccuracy than necessary. As an example, consider a user who wishes to protect a browsing history composed mainly of anime and manga pages. A typical approach to protect this profile would be making it more common, helping the user go unnoticed, in a k -anonymity sense [81, 88]. Nevertheless, without having an idea of how the other profiles are, users cannot adjust the level of perturbation, which may vary significantly depending on whether the profile in question is handed over to an advertising company operating in Japan or in Brazil.

On the other hand, and more importantly, local perturbation ignores the fact that privacy protection is *co-utile* [36], meaning that the best strategy to protect one's privacy is to help others in protecting theirs. The notion of co-utility is rather intuitive when applied to the anonymization of browsing data. In the example above where a user tries to blend their profile into a group of affine peers, none of the members of this group would like any of them to be re-identified; if that happened, their own profiles would be more

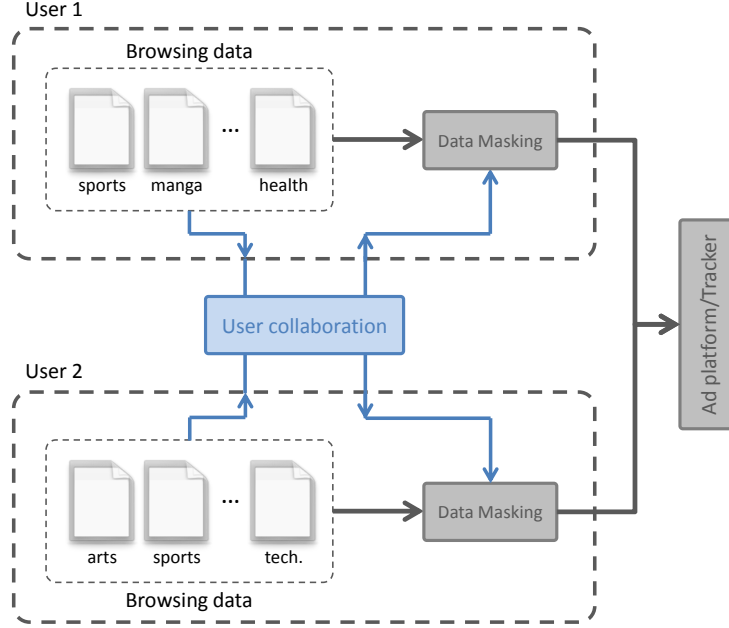


Figure 3: Conceptual depiction highlighting the differences between local and collaborative masking. The data sharing between two users permits them to make an informed decision regarding the level of perturbation.

easily re-identified. In a k -anonymity-like model, for example, this can be straightforwardly quantified: the re-identification of a member within a group of size k would increase the probability of re-identification of the remaining members by $k/(k - 1)$. In a nutshell, what this suggests is that users will not only be interested in guarding their own privacy, but also they will be willing to assist others in preserving theirs. This is the principle behind the notion of co-utility.

Our work relies on the fact that privacy protection is co-utile, which ensures that users will be disposed to collaborate with each other to facilitate the protection of the rest of the users involved; this is an aspect that local masking does not contemplate. We adopt this co-utile collaborative approach by enabling users to share their browsing profiles, and by providing them with data-perturbative methods that can be executed on their side. As we shall describe in the coming sections, our scheme will permit users to protect their data more effectively, by making local but informed decisions about the level of perturbation to be applied. Fig. 3 illustrates, at a conceptual level, the differences between local masking and the scheme proposed in this work based on user collaboration.

3.3. Architecture Overview

Next, we introduce the components of an architecture that supports the tracking model described in Secs. 3.1 and 3.2. These components are a data broker and a Web-browser plug-in that is installed on each user’s machine.

In those two subsections, we assumed that users are willing to hand their browsing profiles over to ad platforms and trackers. Nevertheless, this is only if users can control and monetize the disclosure of such private data. Because having a large number of users negotiating said disclosure directly with ad companies is not practical, this work contemplates a dedicated entity acting as an intermediary between users and companies. We call this entity *user-data broker* and, in essence, it may be regarded as a *browsing data marketplace*. We would like to note that the consideration of such an intermediary entity is in line with several recent data monetization projects like DataCoup, DataWallet and CitizenMe. Although not implemented yet, the idea behind these projects is to allow users to sell their social networks data to advertisers through an intermediary broker.

In the envisioned marketplace, users agree to be tracked by all or some of the ad platforms and trackers registered with the broker; Sec. 4 will describe this subscription system and other more practical aspects.

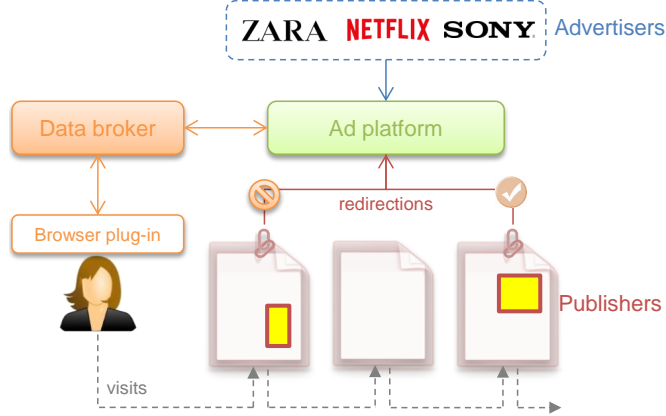


Figure 4: We contemplate a user-data broker and a Web-browser plug-in to support the online tracking model introduced in Sec. 3.1. On the one hand, the broker allows users to negotiate the extent of tracking with ad platforms and trackers. On the other hand, the plug-in is responsible for controlling the access, by these companies, to a user’s browsing profile. The control is done by blocking or allowing redirections to the ad platform or tracker.

However, depending on the compensation given to users, the tracked profiles will resemble their actual browsing interests or not. By allowing tracking even in the absence of compensation or subscription, our approach supports the Internet’s current business model in which content is paid from advertisement revenue. As we shall describe in greater detail in the next subsections and Sec. 4, profiles will be perturbed in real-time (i.e., as users browse the Web) by selectively blocking or allowing tracking on each visited page. In doing so, all forms of advertising —except profile-based ads— will be allowed even when no compensation is provided. Nevertheless, when there is no reward, tracking and advertising will be partly disabled, and the remaining tracked pages will not reflect real browsing habits and hence tracking will be useless for ad personalization. In other words, the conversion rates of interest-based advertising will decrease to those of untargeted ads.

As an intermediary entity, the user-data broker will therefore inform users about the compensation offered by ad companies. Users, in turn, will not only be able to accept or decline such offers, but also they will have the possibility to notify these companies of their own rates for allowing unperturbed tracking. We contemplate that these rates are specified on a per-category basis, and vary over time depending on users’ needs and perceptions on privacy, as well as ad platforms’ requirements to satisfy the advertisers’ demand. However, as in any competitive market, the unit price will be determined on the basis of supply and demand.

Our architecture relies on a Web-browser extension that is installed on the user’s computer and that, together with the user-data broker, enables users to negotiate the scope of tracking. The main functionality of the plug-in, however, is the collaborative masking of the user’s browsing profile.

A crucial aspect of this collaborative masking is the sharing of profile data among users. In this work, we consider any standard zero-knowledge protocol [43] for users to authenticate anonymously with the broker, and to upload their profiles to a shared repository. Also, through this protocol a user will learn about the other users’ profiles, and in particular about the distribution of this information. In this process of data sharing, we shall make two assumptions. First, that user profiles are not identifying information per se; in the terminology of statistical disclosure control (SDC) [48], profiles might be regarded as *key attributes*, meaning information that, only in combination, may be linked with external data to reidentify a user. Secondly, we shall assume that any user participating in this data exchange cannot link an uploaded profile to any other information about the user behind it, which might potentially reidentify them. Our assumptions are in part justified by the user-profile model described later in Sec. 3.4, which relies on a broad categorization of the clickstream —and not on the particular Web pages visited— to represent browsing interests. Sec. 4 will elaborate further on these aspects of data sharing as well as the adversary model considered in this work.

In addition to providing collaborative masking services, through this plug-in ad platforms and trackers will be able to expand their tracking capabilities much beyond current tracking technologies permit; obviously, this will be provided that a tracker is willing to accept users’ rates for full tracking. We believe

this is a great incentive for an ad platform or tracker to be part of the proposed tracking and advertising model: except for Google’s ad platforms, the majority of trackers have a partial view of users’ browsing interests [69, 68, 74]. Consequently, the incorporation of this tool will allow smaller, less-resourced companies to boost their tracking capabilities, and more importantly, these enhanced capabilities will be available with users’ consent.

The proposed architecture will therefore lead to a “win-win” situation for users and trackers. On the one hand, users will be able to exert fine-grained control over Web tracking and will have a real motivation for revealing their browsing profiles. And on the other hand, trackers will have access to complete browsing data through a negotiation process that respects users’ privacy preferences. At a high level, this new online advertising paradigm may contribute to address the current “ad blocking wars”. Fig. 4 illustrates the two components of the architecture described in this subsection.

3.4. Privacy Model and Disclosure Strategy

In the previous subsection we explained that, in order for users to make an informed decision on the level of perturbation, our architecture facilitates the sharing of browsing data. However, to really give users the option to find a “good” trade-off between privacy on the one hand, and on the other the economic reward for tracking, we require translating said perturbation to privacy. That is, we need establishing of a privacy model and a quantitative measure of the privacy of browsing profiles.

In this section, we focus on these two aspects and the relationship between privacy and economic reward. More specifically, we define our models of browsing profile and privacy in Secs. 3.4.1 and 3.4.2, respectively. Afterwards, Sec. 3.4.3 proposes a profile-perturbation technique and a generic measure of privacy risk, and Sec. 3.4.4 presents a formulation of the trade-off between privacy and money.

3.4.1. User-Profile Model

In line with previous works in the literature [90, 77, 58, 93, 28, 74], we model the Web pages browsed by a user as a sequence of random variables (r.v.’s) taking on values in a common finite alphabet of interest categories or topics, in particular the set $\{1, \dots, n\}$ for some integer $n \geq 2$. In our mathematical model, we assume these r.v.’s are independent and identically distributed. This assumption allows us to represent the Web-browsing profile of a user by means of the probability mass function (PMF) according to which such r.v.’s are distributed. Conceptually, we may interpret a profile as a histogram of relative frequencies of visited pages within that set of categories.

We organize this set of categories hierarchically, so that the lower-level interests are more specific and sensitive than those at the higher levels. This organization allows us to use different representations of a same profile, depending on the hierarchical level assumed. In this work, we consider two possible representations of a profile. One version modeled across bottom-level categories, to be disclosed to ad companies, and another version represented across higher-level categories, to be shared among users of our pay-per-tracking service. Accordingly, we shall employ n to refer to the number of bottom-level categories, and define $m \ll n$ as the number of categories of the hierarchical level the user is happy to share with other users.

As we shall describe later on in Sec. 4.1, the availability of this coarser version will permit us to justify, in part, the sensitivity of the sharing of browsing data among users. In particular, we shall assume that their privacy is not compromised by observing this profile in isolation, that is, without any additional information about their browsing activity and related network and browser parameters. We acknowledge, however, that the use of different hierarchical levels for ad platforms and for internal distribution among users, may come at the expense of inaccuracy in determining the desired level of masking. Fig. 5 shows an example of browsing profile with $n = 12$ subcategories and $m = 3$ categories.

Having said that, next we shall assume, unless otherwise stated, that profiles are modeled across n bottom-level categories. Recall that when a user registers with the data broker, they accept unveiling some pieces of their browsing profile, or slightly perturbed versions of it. The user may consider, for example, being tracked on a fraction of their visits to non-sensitive categories, and may decide blocking all tracking on sensitive pages. Whatever the decision is, the upshot is that the ad platform or tracker following these visits observes a perturbed version of the genuine profile, which does not reflect the actual, accurate browsing




Figure 5: We model user profiles as normalized histograms of pages visited across a set of topic categories. The profile represented with $m = 3$ high-level categories is the one users agree to share among themselves to decide the level of perturbation, whereas the profile modeled with $n = 12$ bottom-level categories is the one users perturb and hand over to ad companies.

interests of the user and which therefore may not be useful for ad personalization. In this work, we shall refer to these two profiles as the *actual user profile* and the *apparent user profile*, and we shall denote them by $q = (q_1, \dots, q_n)$ and $t = (t_1, \dots, t_n)$.

3.4.2. Privacy Models

To generate a perturbed version of this actual profile, users must first specify the privacy objective they want to achieve by masking this information. In the literature of information privacy, this objective is inextricably linked to the concrete assumptions about the attacker against which a user wants to protect⁹. For example, the original k -anonymity model, which is the requirement that each tuple of key-attribute values be shared by at least k records in a database, is defined for an adversary that aims at reidentifying the respondents behind those records [81, 88].

In this work, we contemplate two privacy objectives for users, which may also be interpreted from this attacker perspective.

- On the one hand, we assume a *profile-density* or k -anonymity-like model, in which a user wishes to make their profile more common, trying to hide it in the crowd.
- On the other hand, we consider a *classification* model in which the user does not want to be identified as a member of a given group of users.

In terms of an adversary model, the former objective could be defined under the assumption that the ad platform or tracker aims at targeting peculiar users, that is, users who deviate from a typical browsing behavior. The latter model, on the other hand, could fit with an adversary who wishes to label a user as belonging to a particular group. In either case, the ultimate aim of the ad platform or tracker could be from price discrimination to social sorting [60].

The selection of either privacy model entails choosing a reference, *initial profile* p the user wishes to impersonate when no money is offered for tracking them. For example, in the k -anonymity-like model, a user might want to exhibit very common interests and p might therefore be the average profile of the population. In the classification model, the user might be comfortable with showing the profile of a less-sensitive group.

⁹This is known as the *adversary model*. The importance of this model lies in the fact that the level of privacy provided is measured with respect to it.

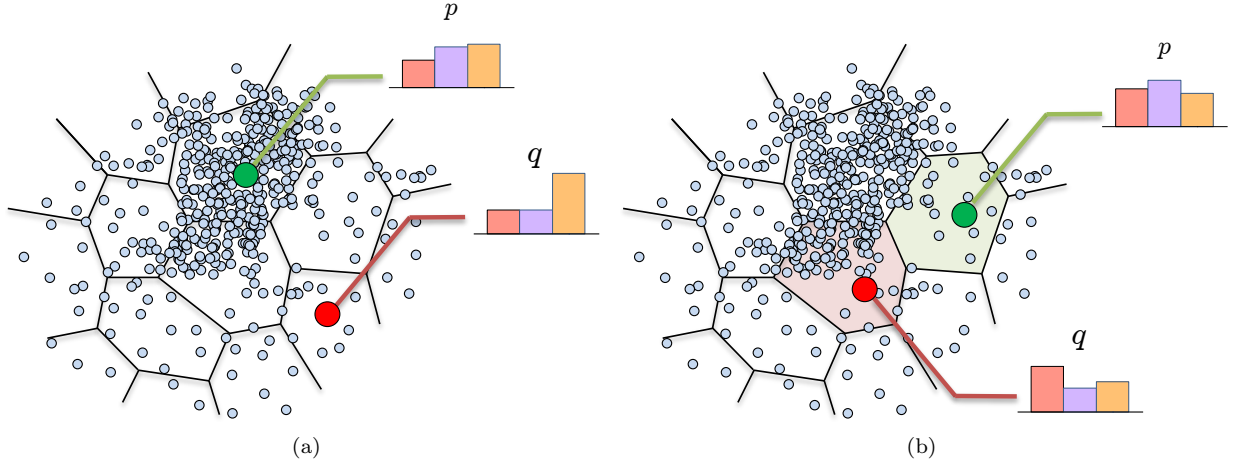


Figure 6: Profile-density and classification privacy models. We depict an example of distribution of browsing profiles for $m = 3$ high-level categories. This distribution is available to all users of our pay-per-tracking service and allows them to make a decision on how to perturb their profiles, depending on the privacy model assumed. Intuitively, in the profile-density model, the perturbation is in direction of moving the actual profile q to a region where there is more density of users (a). In the classification model, on the other hand, a user may want their distorted profile to belong to a cluster where they feel more comfortable (b).

The consideration of such initial profile is crucial for the support of the Internet current business model. As explained at the beginning of Sec. 3.1, the proposed solution is not against tracking and advertising, but only against the fact that users do not have control over the personal information that is collected and exploited to direct them ads. By allowing a profile p —instead of cutting off all tracking—in the absence of compensation or subscription, users can still receive ads, thereby supporting publishers. However, these ads cannot be personalized to their browsing interests. As we shall describe later in Sec. 4, ad platforms and trackers will also experience a reduction in the number of tracked pages, as a consequence of this profile masking.

Fig. 6 illustrates these ideas by means of a simple but insightful example. The figure in question shows a distribution of profiles in the probability simplex, in the case when profiles are modeled across $m = 3$ high-level categories of interest, e.g., education, health and sports. In this figure, we represent the actual and the initial profiles of two particular users, exemplifying the two privacy models considered in this work.

3.4.3. Profile-Disclosure Technique and Privacy Function

In this section, we propose a profile-perturbation technique suitable for the advertising and privacy models described in the previous subsections. The proposed technique operates between these two extreme cases. When there is no compensation for tracking, the masked profile coincides with the initial distribution p , and the observation of this information by an ad company does not pose any privacy risk to the user. When the user is offered sufficient reward, however, the actual profile q is fully disclosed and their privacy completely compromised. Clearly, because the level of perturbation decreases with this reward, we may also interpret the proposed technique as a mechanism that *unmasks or discloses* private information.

Our perturbation mechanism reveals the deviation of the user’s initial, false interest to the actual value. In formal terms, we define the *disclosure rate* δ_i as the percentage of disclosure lying on the line segment between p_i and q_i . Concordantly, we define the user’s apparent profile t as the convex combination [32, §2.1]

$$t = (1 - \delta)p + \delta q,$$

where $\delta = (\delta_1, \dots, \delta_n)$ is some *disclosure strategy* specified by the user. The disclosure mechanism may be interpreted intuitively as a roller blind. The starting position $\delta = 0$ corresponds to leaving the roller in the value p , that is, $t = p$. Depending on whether $q_i < p_i$ or $q_i > p_i$, a positive δ may translate into lowering or raising the roller respectively. Fig. 7 illustrates this effect for a uniform initial profile, that is, $p_i = 1/n$ for all $i = 1, \dots, n$.




Figure 7: Disclosure technique for browsing profiles. We show the apparent profile t that results from applying a certain disclosure strategy on the actual profile q . The initial profile p corresponds to the uniform distribution. The selected strategy fully reveals the interest of the user in the category 3, which means that ad companies can track them on all pages related to this category. However, as for the categories 1 and 2, the user decides to expose half of their interest value (in the line segment between p and q).

Measure of dissimilarity	References
SSD, MSE	[79, 85]
Euclidean distance	[47]
KL divergence	[78, 75]
mutual information	[94, 42]
cosine distance	[39, 38, 54, 84, 37, 40]
Hamming distance	[73]

Table 1: A great deal of research has been devoted to the investigation of privacy metrics. In this table, we show those proposals where privacy is quantified by means of a measure of dissimilarity between profiles. Any of these functions could be adopted in our work.

In our model, the user therefore must decide a disclosure strategy that shifts the apparent profile from the initial PMF to the actual one. The question that follows naturally is, what is the privacy loss due to this shift, or said otherwise, how do we measure the privacy of the apparent profile?

In this work, we do not contemplate a single, specific privacy metric, nor consider that all users evaluate privacy the same way. Instead, each user is allowed to choose the most appropriate measure for their privacy requirements. In particular, we quantify a user’s *privacy risk* generically as

$$\mathcal{R} = f_P(t, p) = f_P((1 - \delta)p + \delta q, p),$$

where $f_P: (t, p) \mapsto f_P(t, p)$ is a *privacy function* that measures the extent to which the user is discontent when the initial profile is p and the apparent profile is t .

A variety of functions may be chosen to reflect this degree of dissatisfaction a user experiences when moving from p towards q . The suitability and appropriateness of the chosen privacy function, however, will hinge on the user’s own perception regarding privacy and the adversary model assumed. Clearly, depending on whether the tracker aims at classifying users or finding uncommon profiles, \mathcal{R} will represent a risk of classification or uniqueness.

A particularly interesting class of those privacy functions are the *dissimilarity metrics*, which have been extensively used to measure the privacy of user profiles. The intuitive reasoning behind these metrics is that apparent profiles closer to p offer better privacy protection than those closer to q , which is consistent with the two privacy models described in Sec. 3.4.2. Examples of these functions comprise the sum of squared differences (SSD), Kullback-Leibler (KL) divergence [34], mean squared error (MSE), cosine distance and Hamming distance. Table 1 lists some works using dissimilarity functions to evaluate privacy.

3.4.4. Formulation of the Trade-Off between Privacy and Money

Equipped with a measure of the privacy risk incurred by a disclosure strategy, users must also indicate the compensation they expect to receive for such disclosure. Clearly, this is the only way to configure a suitable trade-off between privacy and economic reward. Next, we formalize this trade-off by formulating the problem of choosing a disclosure strategy as a multiobjective optimization problem that takes into account privacy and money.

Let $g_i : [0, 1] \rightarrow \mathbb{R}_+$ be a function mapping a disclosure rate in the category i into an economic reward. For the sake of simplicity and illustration, in this work we shall only consider functions of constant values w_1, \dots, w_n . We shall refer to these values as *weights*¹⁰. Accordingly, for a given economic compensation μ , we define the *privacy-money function* as

$$\mathcal{R}(\mu) = \min_{\substack{\delta \\ \sum_i w_i \delta_i = \mu, \\ \sum_i t_i = 1, \\ 0 \leq \delta_i \leq 1}} f_P(t, p), \quad (1)$$

which characterizes the optimal trade-off between privacy and economic compensation.

Conceptually, the result of this optimization problem is a disclosure strategy δ^* that tells us, for a given amount of money, how to unveil a profile so that the level of privacy is maximized. Intuitively, if f_P is a profile-similarity function, the perturbation is chosen to minimize the discrepancies between the apparent and the initial profiles. Naturally, the minimization must satisfy that the compensation offered is effectively exchanged for browsing information. This is what the condition $\mu = \sum_i \delta_i w_i$ means. The other equality condition, $\sum_i t_i = 1$, merely reflects that the resulting apparent profile must be a probability distribution.

In closing, the problem (1) gives a perturbation rule that not only assists users in protecting their privacy, but also gives them the option to find the optimal exchange of privacy for money.

4. User-Side Architecture and some Practical Considerations

In this section, we shall examine how this perturbation rule can be applied in real practice. Specifically, we shall investigate several implementation aspects of the proposed pay-per-tracking service, and describe in greater detail one of the two elements of the architecture described in Sec. 3.3, namely the Web-browser plug-in.

These more practical details will allow us to show how our pay-per-tracking model may provide ad platforms and trackers with more opportunities to track users and enrich their profiles. However, as we shall see in the coming sections, our approach does not create more opportunities for ad companies to deliver ads than those they already have. Said otherwise, to continue serving ads, ad platforms will still need to follow the current advertising practice of embedding their links into the publishers they want to partner with. Obviously, this process must be done in collaboration with publishers, as we explained in Sec. 2.1.2.

Accordingly, our approach may be seen as a *complement* to the current advertising model, in the sense that it provides ad platforms with more data about a user's clickstream, but it does not permit them to display ads on those pages which they do not partner with.

4.1. Assumptions

This section examines the assumptions made in designing the user-side architecture.

We explained in previous subsections that users of our service regulate the tracking done by ad companies. If trackers accept paying users, the collected information may reflect real browsing behaviors. If these companies do not compensate users, they are still allowed to track a fraction of Web pages, they can serve ads there, but the profiles resulting from this tracking are useless for personalized targeting.

Our first assumption is related to the possibility that a tracker can link a user back to *past* browsing records or other information about their interests, and thus can serve profile-based ads without compensating the user. In an attempt to avoid this, before tracking begins through our pay-per-tracking service, our plug-in takes charge of eliminating any cookie and browsing data stored on a user's machine. However, regardless of whether a tracker pays for tracking or not, if the user logs into an online application supported by this tracker, there will be the potential for this entity to effortlessly reidentify the user and/or associate previous activity with them.

¹⁰We assume that these weights are available to ad companies through the user-data broker.

This is the case of Facebook, Twitter and Google, which are companies with extensive tracking capabilities, and with which a user might be registered. In this situation, it is clear that the tracker—and at the same time service provider—will attempt to update its knowledge about the user with any available data, and exploit this information to deliver profile-targeted ads. The only drawback for this entity is that these ads will be allowed only in a fraction of pages.

The linkage described above, however, is not restricted to log-ins into social-networking applications or e-mail services. In fact, a tracker might try to look up a user (if present) in its own database by inspecting their IP address and Web-browser fingerprints, and/or correlating the visited pages and browsing timing activity.

In this work, we assume that ad platforms and trackers are unable to conduct said linkage *unequivocally*. In particular, users either are seen as new users through the eyes of these companies, or they cannot be distinguished among a sufficiently large number of distinct users. In either case, a non-paying tracker finds it useless any attempt to serve interest-based ads. More specifically, we make the following assumptions:

- We assume that trackers cannot *unambiguously* link previous browsing activity back to a user through their IP address, Web-behavioral patterns and browser fingerprints, which renders tracking without compensation useless for personalized advertising. We believe this is a reasonable assumption in the case of dynamic IP addresses, if the plug-in is able to remove uncommon browsing behaviors (see Sec. 4.3), and considering that browser-fingerprinting techniques may be inaccurate and quite effectively counteracted [55, 56].
- On the other hand, we assume users are disposed to block any possibility of tracking by social networks and related service providers, with which users are registered and they have explicitly or implicitly conveyed personal information. Therefore, since such companies are not allowed to deliver any kind of ad, we prevent them from exploiting said information.

Our second assumption has to do with the sharing of profile data and the sensitivity of such sharing. In Sec. 3.2, we mentioned that authenticated users need to share their browsing profiles to decide the perturbation strategy. This work assumes that the observation of a profile by *a user* does not pose a privacy risk to the user behind that profile. This assumption is based on the following considerations. First, users—in contrast to ad platforms and trackers—are unaware of the particular Web pages visited, the user’s location (through their IP address), navigation timing, and other information related to their Web browser and operating system. Secondly, the profiles shared among users themselves are modeled across a reduced set of $m \ll n$ coarse and general categories of interest. And thirdly, the sharing of those profiles is supposed to be done, in practice, among a sufficiently large number of users spanning a vast geographic area, for example, a country. In a nutshell, we assume that both the set of categories and origin of those profiles have a sufficiently large level of granularity so that the probability of reidentifying a user only through their coarse profile is negligible.

Our third and last supposition concerns the main component of the user-side architecture, which assumes that the browsing profile remains stable over a long period of time. We suppose that this steady-state condition is attained after a training phase, once the user has visited a sufficiently large number of pages. This assumption is in line with the so-called long-term profiles which, in contrast to the short-term profiles, reflect interests that are not subject to frequent changes [44]. We acknowledge, nevertheless, that a practical implementation of this architecture should contemplate that the user’s browsing interests might vary significantly with time.

4.2. Subscription to Pay-Per-Tracking

We conceive our per-pay-tracking approach as a service ad platforms and trackers can subscribe to. Subscriptions are planned to be free of charge and the domains of the subscribed ad companies would be available to users through the broker. Before proceeding, we would like to emphasize that a subscribed tracker does not necessarily mean a paying tracker.

As described in the previous section, the idea behind profile perturbation is showing an initial profile p when no money is offered, and revealing the user’s actual profile q when $\mu = \mu_{\max} = \sum_i w_i$. For

an intermediate value $\mu \in (0, \mu_{\max})$, the solution δ^* to the optimization problem (1) gives an apparent, combined profile $t^* = (1 - \delta)p + \delta q$.

As we shall explain in Sec. 4.3.4, our plug-in will have to emulate this browsing behavior t^* while the user navigates the Web. To this end, it will block or allow the third-party network requests that may be initiated when the browser loads a page. However, since tracking is fortunately not ubiquitous, in some cases our plug-in will need to inform trackers of the visits to pages where they are not present. We consider sending a *notification* with the visited page in such cases. In practice, this notification can be seen as if our plug-in executed a cookie-matching protocol [41], through which it redirects the browser to the ad platform or tracker domain in question, including the user’s cookie and the page as parameters in the URL.

When visiting a page, our plug-in therefore will block or allow a tracker’s redirection (assuming its link was previously embedded by the publisher), or send a visit notification if the tracker is not present. Our masking algorithm will mimic the browsing pattern t^* by blocking or allowing tracking this way. Nevertheless, the choice of allowing an existing third-party network request or sending a visit notification will depend on whether the ad platform or tracker is subscribed to our pay-per-tracking service. We contemplate the following three cases:

- **Case 1.** A subscribed tracker is not found in the page.
- **Case 2.** A subscribed tracker is found in the page.
- **Case 3.** A non-subscribed tracker is found in the page.

In the former case, the plug-in will suggest sending a visit notification to the tracker. However, it will be up to the masking module (described later in Sec.4.3.4) to decide if it is finally submitted or not depending on μ and other parameters.

In the latter two cases, on the contrary, our plug-in will leverage the third-party network request, and the possibility that there is an iframe¹¹ available to the ad platform. But again it will be the masking algorithm that determines if the request is blocked or allowed. It is important to stress that, only in cases 2 and 3, the ad platform may be in a position to deliver an ad. Put differently, when our plug-in sends a visit notification to a subscriber, this does not enable the ad platform to display an ad on the notified page.

Another key point is the fact that our plug-in may allow third-party connections from non-subscribed trackers. We permit this possibility because our solution aims at preserving the current Internet business model. Although user privacy is the priority, our service is devised so as to cause the least possible impact on the Web economy. We shall see it later in Sec.4.3.4, where our perturbative algorithm, despite the necessary blockage of a non-negligible number of visits, is designed to minimize such effect.

On the other hand, because allowing tracking on those pages where ad companies are not present (case 1) comes at the cost of traffic overhead, users will be the ones to choose the list of subscribers they are willing to support. However, a related question arises in this case 1. Since it does not distinguish whether the subscribed tracker is a paying tracker or not, why should a user send visit notifications to the non-paying ones if this comes at the expense of traffic?

The answer is a simple one: these notifications may allow users to advertise themselves as potential clients for ad platforms and trackers. We know that a subscribed, non-paying tracker will observe a profile $t = p$, which does not help for ad personalization. Besides, visit notifications will not increase a tracker’s opportunities for serving ads. In fact, these opportunities may be reduced due to the mimicking effect (see Sec.4.3.4). Consequently, compared to a non-subscribed tracker, an ad platform or tracker which is subscribed to our service but which does not compensate a user for tracking will *only* benefit from a better estimation of the number of pages they browse. Obviously, this gives the ad platform an idea of the reliability of the estimation of the user’s interests and the effectiveness of interest-based ads if it ends up paying this user. In short, sending visit notifications to subscribed but not paying ad companies helps them decide whether to invest or not in a user, which in turn is useful for the user themselves. We shall elaborate more on this in Sec.4.3.4.

¹¹Iframes are typically used to serve ads from third-party domains.

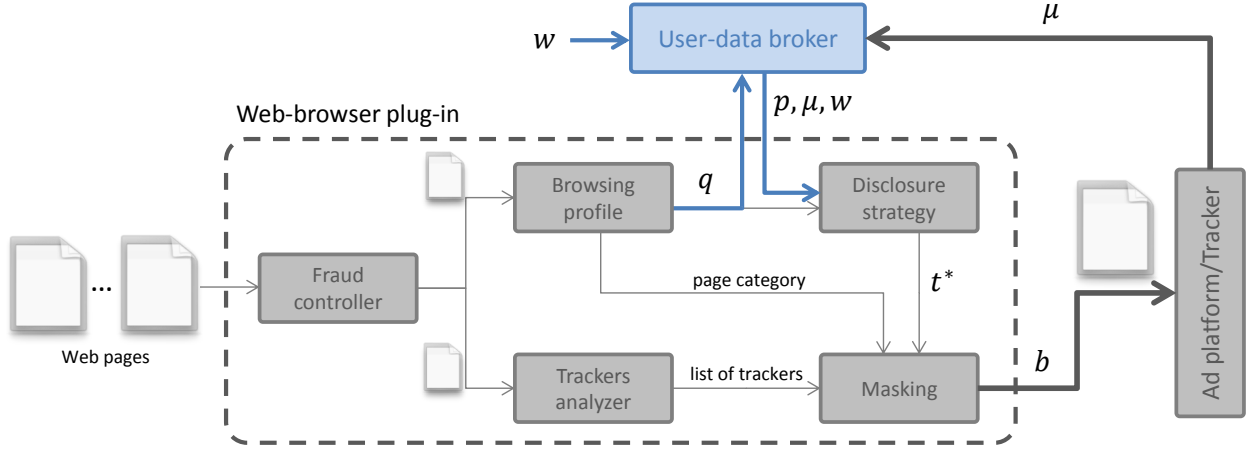


Figure 8: Internal components of the proposed user-side architecture.

Finally, we would like to note the need of storing a user profile for each non-subscribed tracker. As we require emulating p and the plug-in can only manipulate (i.e., block) to this end the present third-party network requests, the masking algorithm will need first to estimate the user profile in their hands, tracked through their partnering publishers. Recall from Sec. 2.1 that a user’s clickstream is precisely the sequence of pages browsed by a user and tracked by an ad platform or tracker. Throughout this section, we shall refer to it as the *observed clickstream* to distinguish it from the *actual clickstream*, that is, the sequence of *all* pages visited by a user, from which q is built. We shall denote by q' the profile that results from this observed clickstream. Clearly, depending on the ability of the ad platform or tracker to track users throughout the Web, the profile q' will resemble, to a greater or lesser extent, the profile q .

4.3. Modules

The user-side architecture described in this subsection aims at protecting user privacy by means of collaborative masking methods. Unlike the traditional methods introduced in Sec. 3.2, privacy protection is carried out locally albeit informedly. The proposed plug-in optimizes the masking of profiles and perturbs browsing information in real time, i.e., as the user browses the Web. The optimal disclosure strategy determines in which pages we should allow or block tracking so that privacy is maximized for a given economic reward. The fundamental purpose of this plug-in, and therefore of this architecture, is to help users regain control over tracking.

In this section, we shall explore the main functionalities of this tool, and specify the conceptual design and fundamental operational structure of a practical implementation. We would like to stress that the description provided in this section does not pretend to serve as an exhaustive guide for programming such tool.

Fig. 8 depicts the proposed architecture, which consists of a number of components each of them performing a specific task.

4.3.1. Trackers Analyzer

This module is responsible for finding out the ad platforms and trackers that might potentially track a user’s visit to a Web page. With this aim, it first captures all third-party network requests which may be initiated on each visited page. Since trackers are not the only domains which may be involved in these requests¹², the module will need to filter them. To identify the trackers, it may rely on the lists used by ad blockers and anti-trackers to this end. As mentioned in Sec. 2.2, these tools typically keep a list of domains serving ads and/or tracking users, which is publicly available. Some examples include EasyList [13] and Ghostery’s [9]

¹²Quite often publishers serve content from other domains, for example, from distinct domains owned by the publisher or content delivery networks.

list. Finally, after this filtering is applied, the module sends the tracker domains found in the page to the module *masking*.

4.3.2. Browsing Profile

It is in charge of estimating the two versions of the user’s actual profile, one with m general categories, and another with $n \gg m$ specific and more sensitive categories.

The module is equipped with a local Web-page categorizer to classify a user’s visits into those topic categories. Examples of categorizers that could be employed with this purpose include the ones available at the Firefox Interest Dashboard plug-in [19] and MyTrackingChoices [25, 26].

To estimate q from these categorized pages, a practical implementation of the architecture could rely on maximum-likelihood (ML) estimation [83] or the Laplace’s rule of succession [83]. As mentioned in Sec. 4.1, our plug-in assumes that, when estimating this profile, the relative frequencies of activity are sufficiently stable after a significant number of Web visits. We shall suppose a training period of 2 weeks, which, according to some surveys, may be the time period required by ad platforms to model long-term interests [72]. Obviously, all tracking will be blocked during this training phase.

Evidently, the same methodology and training phase applies to the estimation of the user profiles available at ad platforms and trackers. The difference between q and q' is in that the former is built locally on the basis of *all* pages visited by a user, whereas the latter is computed from the tracked pages, that is, from the pages ad platforms and trackers have embedded a link to.

Last but not least, in Sec. 4.1 we commented that a tracker might try to link a user back to own external records based on their browsing patterns. To mitigate this type of linkages, during the training and tracking phases the module will check the popularity of the visited pages with Alexa’s top 100 000 sites, and block all trackers on unfrequented sites.

4.3.3. Disclosure Strategy

This module is responsible for computing the solution to the optimization problem (1). To this end, it is provided with the actual user profile q , namely, with the version across the n bottom-level categories.

Through the user-data broker, the user will learn the distribution of actual profiles of the other users and select, accordingly, an initial distribution p they are comfortable to show when no money is offered to them. However, because the shared profiles are modeled across m high-level categories, the module at hand will need to convert p into an n -dimensional representation. For this, we consider two options. In the former approach, the module fills the bottom-level categories with a uniform value equal to the corresponding high-level interest. In the latter approach, it is the user who selects the particular values of these bottom-level interests, bearing in mind the high-level interest.

The disclosure-strategy module will also be informed of the economic compensation μ offered by the tracker (through the broker), and will require the user to specify a privacy function f_P and a tuple of weights $w = (w_1, \dots, w_n)$. With all this information, the module will rely on optimization libraries to compute the optimal disclosure strategy δ^* and the optimal apparent profile t^* . For the sake of brevity, the description of this module is restricted to the case of a single tracker. We must take into account, though, that a user could choose different values of the parameters p, μ, f_P, w for each tracker, and therefore δ^* and t^* could be specific to each of them.

The obtained apparent profile t^* is the centerpiece of the plug-in described in this section, and the information that is sent to the *masking* module. This profile represents the browsing activity that the tracker in question should observe. In particular, the component t_i^* is the percentage of pages belonging to the category i that the plug-in should allow this tracker to track.

Finally, we would like to emphasize that the particular privacy function chosen will be the key factor determining the implementation of this module. Some examples of open-source, efficient optimization solvers include Coin-OR Interior Point OPTimizer (IPOPT) [4, 92], Limited Memory Broyden-Fletcher-Goldfarb-Shanno (LBFGS) [97] and NLOpt [51].

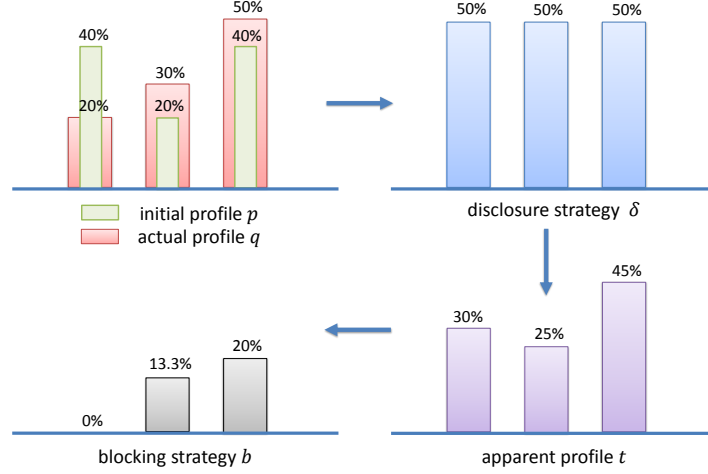


Figure 9: Example of how a disclosure strategy is transformed into a blocking rule. The example shows an actual user profile $q = (0.20, 0.30, 0.50)$ and a reference profile $p = (0.40, 0.20, 0.40)$. We suppose that, for a given μ , the solution to the optimization problem is a disclosure strategy $\delta = (0.50, 0.50, 0.50)$, from which we obtain the apparent user profile $t = (0.30, 0.25, 0.45)$. To emulate a browsing behavior like this, we compute a blocking strategy b that tells us how to turn on and off tracking so that the profile observed by the tracker resembles t . In this case, $\beta_{\min} \simeq 0.33$ and $b \simeq (0, 0.13, 0.20)$, which means that the plug-in will block approximately 43.3% and 40% of the pages corresponding to the categories 2 and 3, respectively.

4.3.4. Masking

As with the previous module, the description of this masking component is restricted to the case of one tracker or ad platform. We first consider this tracker or ad platform is subscribed to our per-pay-tracking service. Afterwards, we examine the simpler case of non-subscribers.

The aim of this block is to translate the optimal disclosure strategy δ^* into network-level actions that make trackers believe the user has a profile t^* . The difficulty of this task lies in that the perturbation must be applied in real time.

The module at hand implements a *masking algorithm* that is able to emulate the browsing pattern dictated by t^* . The algorithm mimics this profile by enabling and disabling tracking on a per-page basis, depending on a series of parameters. The effect of turning on and off tracking may be interpreted, intuitively, as if we kept or subtracted the pages that compose the actual clickstream. This intuition is formalized next.

Let $\beta \in [0, 1]$ be a *blocking rate*, that is, the ratio of blocked pages to total browsed pages the user is willing to subtract to reproduce t^* . Let $b = (b_1, \dots, b_n)$ be a *blocking strategy*, with b_i denoting the percentage of pages that should be blocked in the category i . Formally, we model the effect of alternating between blocking and allowing by means of the distribution $\hat{t} = \frac{q-b}{1-\beta}$. Conceptually, this distribution may be interpreted as the result of, on the one hand, the suppression or blockage of certain pages from the actual clickstream, that is, $q - b$, and on the other, the subsequent normalization by $\frac{1}{1-\beta}$ so that \hat{t} can be a PMF.

Accordingly, our masking algorithm aims at finding the tuple b such that

$$\hat{t} = \frac{q - b}{1 - \beta} = t^*.$$

Obviously, the chosen strategy must satisfy that $0 \leq b_i \leq q_i$ and that $\sum_i b_i = \beta$. Since b is nonnegative, it is immediate to check that the only condition consistent with $q_i = 0$ for some $i = 1, \dots, n$, is that $t_i^* = 0$. In words, this means that a blocking strategy cannot induce a positive interest in those categories where the user has no interest. In practical terms, this implies that the user will be prevented from choosing an initial profile p with positive components where their actual interests are zero.

On the other hand, it is easy to verify the existence of a *critical* blocking rate

$$\beta_{\text{crit}} = 1 - \min_i \frac{q_i}{t_i^*} \quad (2)$$

beyond which the matching $\hat{t} = t^*$ is achieved. Since the proposed advertising architecture aims at causing the least possible impact on the Web economic model, β is chosen to be this minimum rate by default. The

Algorithm 1: Masking algorithm. We assume the same reward μ and tuple w for all subscribers.

Input: p, q, μ, w ; the profile q' for each non-subscribed tracker.
Output: A decision on whether each tracker must be blocked, allowed or notified.

```

1 let  $\mathcal{A}$  be the set of subscribers available at the plug-in.
2 while 1 do
3   let  $i$  be the category of the visited page.
4   let  $\mathcal{B}$  be the set of trackers on that page.
5   for all tracker in  $\mathcal{A}$  do
6     Use  $p, q, \mu, w$  to compute  $\delta^*$  and  $t^*$  from (1).
7     Compute  $\beta_{\text{crit}} = 1 - \min_i \frac{q_i}{t_i^*}$  and  $b = q - t^* \beta_{\text{crit}}$ .
8     if the tracker is in  $\mathcal{B}$  then
9       with probability  $1 - \frac{b_i}{q_i}$ , allow it.
10      with probability  $\frac{b_i}{q_i}$ , block it.
11    else
12      with probability  $1 - \frac{b_i}{q_i}$ , notify it.
13    end
14  for all tracker in  $\mathcal{A} \cap \mathcal{B}$  do
15    Compute  $\beta'_{\text{crit}} = 1 - \min_i \frac{q'_i}{p_i}$  and  $b = q' - p \beta'_{\text{crit}}$ .
16    with probability  $1 - \frac{b_i}{q'_i}$ , allow it.
17    with probability  $\frac{b_i}{q'_i}$ , block it.
18  end
19 end

```

complement of this rate, $1 - \beta_{\text{crit}}$, is precisely the fraction of Web traffic a tracker is allowed to access. The blocking strategy for this $\beta = \beta_{\text{min}}$ yields

$$b(\mu) = q - t^* \min_i \frac{q_i}{t_i^*}, \quad (3)$$

where we emphasize the fact that this strategy (and also β_{crit}) depends on μ . Note that when $\mu = \mu_{\text{max}}$, it follows that $t^* = q$, $b_i = 0$ for all i , and hence all pages are consistently allowed to be tracked.

Based on the tuple b , the masking algorithm may then be straightforwardly implemented. When the user visits a page, the masking module waits for the browsing profile module to send the category corresponding to that page. Let i be the index of this category. The module computes the component b_i and proceeds as follows. If the tracker or ad platform is present on that page, with probability b_i/q_i the plug-in prevents it from making a third-party network request, whereas with probability $1 - b_i/q_i$, this request is allowed. If the tracker is not present on the page, with probability $1 - b_i/q_i$ the plug-in sends a notification of this visit to the tracker. Fig. 9 shows an example of a blocking strategy and how we emulate an apparent profile by blocking and allowing the visits of the actual one.

Having shown the case of subscribers, now we briefly examine what happens if the tracker is not subscribed to the proposed service. The main difference here is that the enabling and disabling of tracking is not applied to the actual clickstream but to the observed one. Also, because $\mu = 0$, our plug-in need not compute the solution to the problem (1). In this case, it suffices to replace q with q' , and t^* with p , in the expressions (2) and (3).

In light of the algorithm described above, it is important to emphasize the difference—that we anticipated in Sec.4.2—between a subscribed, non-paying tracker and a non-subscribed tracker, when they attempt to estimate a user’s activity. In both cases, the estimation of the number of pages browsed by the user is done from the profile p . However, in the former case this profile results from blocking certain pages of the actual clickstream, whereas in the latter case, p results from blocking certain pages of the observed clickstream. It

Table 2: Description of the main variables used in our notation. In gray, we show the parameters introduced in Sec. 4.

Symbol	Description
n, m	number of bottom-level and top-level categories into which a Web page is respectively classified
q	the <i>actual</i> user profile is the genuine profile of interests, built from the actual clickstream
p	the <i>initial</i> user profile is the profile the user is happy to impersonate when no money is offered to track them
μ	an economic reward or <i>compensation</i> offered by a tracker
w	the price fixed by a user for fully unfolding their browsing interests in each category
δ	a <i>disclosure strategy</i> is an n -tuple with percentages of interest disclosure per category
t	the <i>apparent</i> user profile is the perturbed profile resulting from the application of a disclosure strategy
f_P	the <i>privacy function</i> measures the degree of dissatisfaction the user shows when the initial profile is p and the apparent profile is t
$\mathcal{R}(\mu)$	function modeling the privacy-money trade-off
q'	the profile of interests that a tracker builds from a user's visits to publishers partnering with it
\hat{t}	activity profile that attempts to reproduce the browsing behavior described by t through the blockage of certain page visits
β	a <i>blocking rate</i> is the percentage of pages a user is disposed to block to mimic the optimal t^*
β_{crit}	minimum blocking rate to attain $t = t^*$
b	a <i>blocking strategy</i> is an n -tuple that contains the percentage of pages that should be blocked per category from q or q' , to emulate t^*

is not possible to know if

$$\beta'_{\text{crit}} = 1 - \min_i \frac{q'_i}{p_i} < \beta_{\text{crit}} = 1 - \min_i \frac{q_i}{p_i}.$$

But because a user's activity is expected to be greater than that observed by a tracker, although $\beta'_{\text{crit}} > \beta_{\text{crit}}$, the estimated number of page visits by a non-subscriber will very likely be smaller than a subscriber's estimate. In short, a subscriber will in principle have a better estimate. That said, we interestingly note that, if this is not case, the non-subscriber's estimate provides no benefit to this tracker if it afterwards registers with our service. The reason is that we assumed in Sec. 4.1 that trackers cannot unambiguously link previous browsing activity back to a user of our service.

Lastly, we would like to stress that our proposal contemplates any form of tracking by ad companies and trackers. By observing the third-party network requests, our plug-in captures the pages these entities may track through HTTP cookies or *other* more sophisticated methods like Web-browser fingerprinting. By blocking or allowing these network requests, the algorithm described above may control third-party tracking regardless of the tracking method used. The notation used throughout this work is summarized in Table 2.

4.3.5. Fraud Controller

The proposed tracking model allows ad companies to obtain accurate browsing profiles when they are willing to compensate users for this sensitive, private information. The aim of this component is to preclude users from committing fraud by exhibiting behaviors that do not reflect their real browsing interests.

There exist two possible ways for a user to disguise their actual profile. The former is mimicking visits to pages which are not aligned with their preferences. The latter consists in preventing these companies

from tracking visits to pages the user is interested in. We may refer to them as *forgery* and *suppression* disguising techniques, respectively.

Among the two, we may feasibly assume that the forgery technique may only be done manually, that is, by the user themselves. In the last years, the advertising industry has developed a variety of solutions to combat the fraud committed by non-human traffic or bots registering impressions and clicks on ads [50]. Therefore, if the mechanisms to fight against bots-based ad fraud are effective, they are so against the forgery techniques considered here. In other words, our pay-per-tracking service does not require any additional mechanism to prevent forgery-based, profile disguising techniques than those currently available to ad companies. Also, the fact that we model user profiles as normalized histograms makes it useless for a user to randomly visit pages with the aim of increasing the volume of data to sell. The deal in our pay-per-tracking service is not about quantity but accuracy of browsing data: the more trackers pay, the more similar the perturbed and the actual profiles look.

The proposed plug-in, however, does consider the incorporation of mechanisms to counter, or at least mitigate, suppression strategies. In particular, the software may attempt to prevent the use of anti-trackers and anti-blockers, and the connection to anonymization services that may block tracking as well. Currently, most Web-browsers allow detecting the extensions a user is running, which may permit our plug-in to ascertain whether they are also trying to block tracking on their own. In addition, a practical implementation of this plug-in will monitor network traffic to find out if the user is employing other Web browsers where it has not been installed. If any of these circumstances is observed, our plug-in will immediately cancel the user's pay-per-tracking subscription.

5. Evaluation

In this section, we evaluate the impact that our pay-per-tracking approach would have on a variety of aspects, ranging from the tracking capabilities of ad companies and the effectiveness of ad personalization, to the Internet economy and user privacy. With this experimental evaluation, we aim at demonstrating the technical feasibility of our scheme, and the benefits it would bring to both ad companies and users.

5.1. Data Set

Our analysis has been conducted on the basis of the browsing data and ads of 144 users of the Web-browser extension *MyTrackingChoices*¹³. The tool that has allowed us to collect these data was developed within our project *MyRealOnlineChoices*¹⁴ and, fundamentally, enables users to choose the categories of the pages (e.g., adult, religion) where they do not want to be tracked.

Our data set is composed of the following information: the categories declared by users as sensitive, the categories of the visited pages, the trackers present on each of those pages, and the source URLs of all iframes. It is worth noting that these data correspond to users who downloaded *MyTrackingChoices* directly from the Chrome Web Store. Said otherwise, they were not specifically recruited to participate in our experiment.

Our series of experiments were run from April 12 to May 21, 2016, and allowed us to capture 119 026 page visits and 66 477 ads. The data set used in our analysis was preprocessed so as to retain users with a sufficient browsing activity. In particular, we decided to get rid of those users who visited less than 100 pages, since, typically, this is the minimum length of the time window used by ad platforms to build a profile [72, 28, 16]. After this preprocessing, the number of users, pages and ads reduced to 82, 117 125 and 45 710, respectively.

Recall that an important aspect of our pay-per-tracking scheme is the categorization of the pages browsed by users (see Sec. 4.3). In our data set, the pages visited by users are categorized into 32 topic categories. The categorization of the browsed pages has been carried out on the user side by our data collection tool, through an algorithm that is partly inspired by the methodology presented in [52] for classifying non-textual

¹³<https://chrome.google.com/webstore/detail/mytrackingchoices/fmonkjimgifgcgeodhhgbfoncmjclka>

¹⁴<https://myrealonlinechoices.inrialpes.fr>

Table 3: Interest categories.

Index	Category name	Index	Category name	Index	Category name	Index	Category name
1	adult	9	economics	17	hobbies & interests	25	politics
2	agriculture	10	education	18	home	26	real estate
3	animals	11	family & parenting	19	law	27	religion
4	architecture	12	fashion	20	military	28	science
5	arts & entertain.	13	folklore	21	news	29	society
6	automotive	14	food & drink	22	personal finance	30	sports
7	business	15	health & fitness	23	pets	31	tech. & computing
8	careers	16	history	24	philosophy	32	travel

ads into interest categories. Further details about this algorithm can be found in [25, 74]. Table 3 shows the categories of the taxonomy used.

5.2. Privacy Models

We evaluate our pay-per-tracking scheme for the two privacy models described in Sec. 3.4. In particular, we study the case when all users adopt a profile-density model and thus want their initial profiles to show common interests, and the case when these same users opt for a classification approach.

In both scenarios, we assume the KL divergence as privacy function,

$$f_P(t, p) = D(t \| p) \doteq \sum_i t_i \log \frac{t_i}{p_i},$$

where the logarithm is taken to base 2. The KL divergence has been extensively used as a privacy metric [75, 78] and as a classifier in image recognition, machine learning and in information security [70, 76, 49, 45, 71, 63]. Although the KL divergence is not a distance function, because it is neither symmetric nor satisfies the triangle inequality, it does provide a measure of discrepancy between distributions, in the sense that $D(t \| p) \geq 0$, with equality if, and only if, $t = p$.

The consideration of either of the two privacy models assumed in these experiments requires the definition of an initial profile for each user. We proceed by applying Lloyd’s algorithm¹⁵ to group all users of our data set into 10 clusters; this number of clusters has been chosen to achieve a granularity level sufficiently aggregated and thus avoiding groups with few profiles. We define the average profile of the population \bar{p} as the centroid of the largest group. Accordingly, in the profile-density model, all users configure this distribution as the initial profile, while in the classification model, each user is assigned the centroid of the more distant group, in cosine distance [61].

It is important to stress that the assignment of the initial profiles is made bearing in mind that the masking algorithm, described in Sec. 4.3, cannot emulate browsing activity in those categories where a user shows no interest at all. Consequently, we define each user’s initial profile in accordance with their actual distribution. Formally, $p_i = \bar{p}_i$ for all i such that $q_i > 0$, and $p_i = 0$ otherwise. This is valid for the profile-density model. As for the other model, the assignment is made analogously, by replacing \bar{p} with the furthest centroid.

Fig. 10 represents the average profile \bar{p} , which shows a clear bias toward the interest categories “technology & computing” and “science”. Fig. 11 depicts, on the other hand, the number of users per group. The representative profile of each of these groups is shown in Appendix Appendix A. In addition, we plot in Fig. 12 the PMF of the number of active categories, which gives us an idea of how users’ initial profiles will approximate to \bar{p} . As a reference, the cardinality of the support set of this latter profile is 23, while users average 16.3 active categories.

Finally, we would like to emphasize that the clustering and the choice of the initial profiles have been made for $m = n = 32$ categories. That is, in our experiments we consider that profiles, regardless of whether they are shared with other users or sold to trackers, they are all modeled across 32 interest categories.

¹⁵Lloyd’s algorithm [59], which is normally referred to as k -means in the computer science community, is a popular iterated algorithm for grouping data points into a set of k clusters.




Figure 10: Average profile of the population \bar{p} . The interest categories are indexed in Table 3.

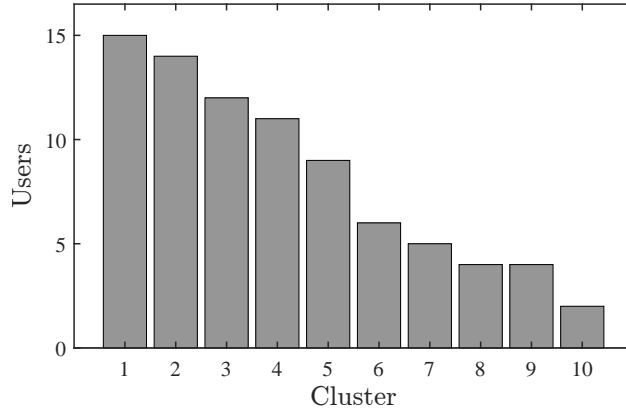


Figure 11: Number of users per cluster. The centroid of the first group, shown in Fig. 10 is defined as the profile \bar{p} .

We note, nevertheless, that a practical implementation of our scheme should choose a number of top-level categories m much smaller than the number of subcategories n , so as to prevent the sharing of very specific interests among the users themselves.

5.3. Transition Scenarios

The evaluation of our pay-per-tracking service is conducted on the basis of three *transition scenarios*. In the scenarios A and B, we consider the case when we move from the current tracking and advertising model, where our pay-per-tracking service is not available yet, to a new context where this service is fully deployed. In both scenarios, we assume all trackers in our dataset are subscribers. In the scenario A, however, trackers are not willing to compensate users for tracking, whereas in the scenario B trackers are disposed to pay users $\mu = \mu_{\max}$ for completely disclosing their profiles. The last transition scenario, the scenario C, contemplates the case when pay-per-tracking is implemented and the subscribers go from no paying for tracking to compensating users for it. Table 4 illustrates the scenarios considered in our experiments.

We would like to emphasize that the former two cases may in fact be regarded as worst-case scenarios. In the scenario A, a tracker goes from watching a whole observed clickstream (and the subsequent profile q'), to observing a profile p that results from blocking certain pages of a user's actual clickstream. This case is certainly worse than the transition scenario that goes from non-subscribed trackers to subscribed, non-paying trackers. Here, the transition would start with a tracker observing p , but this profile would be obtained by blocking certain pages of the observed clickstream. Since the number of pages tracked this way is obviously smaller than those available at the initial state of scenario A, this would give us greater *gains* in the number of tracked pages. Because an entirely analogous argument applies to the scenario B, the results provided in the coming sections for the scenarios A and B may be considered as lower bounds on the enhanced tracking capabilities due to our pay-per-tracking service.

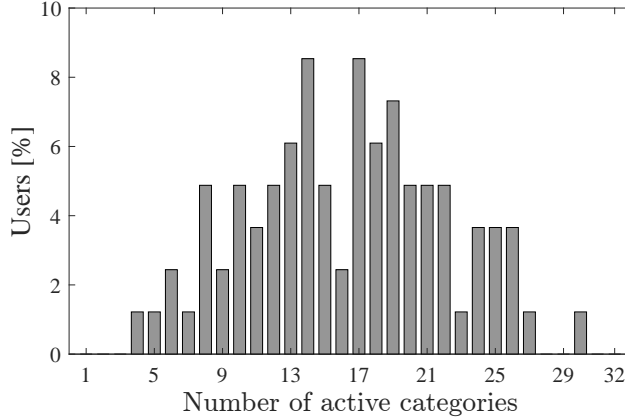


Figure 12: Probability distribution of the number of users' active categories.

Table 4: We analyze the proposed pay-per-tracking model for the three scenario shown here. We assume that, in the new model, all trackers are subscribers.

	Transition	
	From	To
Scenario A	current model	new model, $\mu = 0$
Scenario B	current model	new model, $\mu = \mu_{\max}$
Scenario C	new model, $\mu = 0$	new model, $\mu \in (0, \mu_{\max}]$

5.4. Results

This section is devoted to evaluate several interesting aspects of tracking and advertising for the three scenarios described above.

5.4.1. Scenarios A and B

This subsection analyzes the proposed pay-per-tracking model for the scenarios A and B. Recall that the former considers the transition from the current advertising and tracking model to the new in the case when subscribed trackers do not want pay users for tracking. The scenario B differs from A in that trackers are willing to pay users $\mu = \mu_{\max}$ to learn their complete actual profiles.

In our first series of experiments, we evaluate the extent to which our approach may help ad companies enhance their tracking capabilities.

Fig. 13 shows the relative increase in tracked pages, averaged over all users, for the scenario A and for the 30 most active trackers. In this figure, the trackers are sorted in decreasing order of total tracked pages, and the values of relative increase are represented for the two privacy models specified in Sec. 5.2.

Several conclusions can be drawn from this figure. First and foremost, both privacy models show enormous rates of increase, in the order of hundreds to thousands. The average and maximum values are 581.2% and 2142.8%, respectively. However, we must not deceive ourselves into thinking that this may help trackers improve their ability to *profile* users: the profiles built from these new pages do not reflect real browsing interests in the scenario A. The tremendous benefits reported in terms of tracking are true for all tracker domains, except for `doubleclick.net` and when the privacy model assumed is classification. For this particular tracker and model, we have seen a *reduction* in the number of tracked pages by 4.71%. This is not entirely unexpected since `doubleclick.net` is the tracker with more users (141) and more pages (18347) tracked overall.

Conforming to intuition, we also observe in this figure that those trackers with a smaller presence on the Web are those with greater improvements in tracking capabilities. For instance, while `facebook.com` experiences an increase of 128.8% and 92.0% in tracked pages respectively for the profile-density and classification models, an ad company with limited tracking capabilities like `demdex.net` obtains tracking gains between 1313% and 1700%. On the other hand, we notice that classification yields, in general, worse results. The




Figure 13: Relative increase in the number of tracked pages in the scenario A.




Figure 14: Average percentage of blocked pages when ad platforms and trackers are not subscribed to our pay-per-tracking service.

reason for this is clearly in the larger distance between the actual profile and the initial PMF. This latter distribution is, after all, chosen to be the most distant centroid, which obviously results in higher values of β_{crit} , or blocked pages.

Fig. 13 showed high rates of improvement in terms of tracked pages, in a scenario where trackers are not willing to compensate any single user, and despite the fact that the tracked pages result from a selective blocking of the actual clickstream (see Sec. 4.3). However, we must bear in mind that, when removing pages from this clickstream, our masking algorithm also affects the observed clickstream, that is, the pages that ad platforms directly track and where they may display ads on. In other words, the enhanced —although useless— tracking capabilities of a non-paying subscriber come at the cost of both page and ad blocks.

Clearly, this effect does not only occur to non-paying subscribers but also to non-subscribers. In this latter case, certain pages of the observed clickstream (rather than the actual one) will be blocked to emulate the profile p . In Fig. 14, we show the percentage of blocked pages β'_{crit} averaged over all users, for each non-subscribed tracker. The ad platforms and trackers are sorted in decreasing order of total number of tracked pages, analogously to Fig. 13. Unfortunately, we observe a significant reduction of the observed clickstream, which seems to decrease with trackers' presence on the Web. On average, a tracker would be blocked on

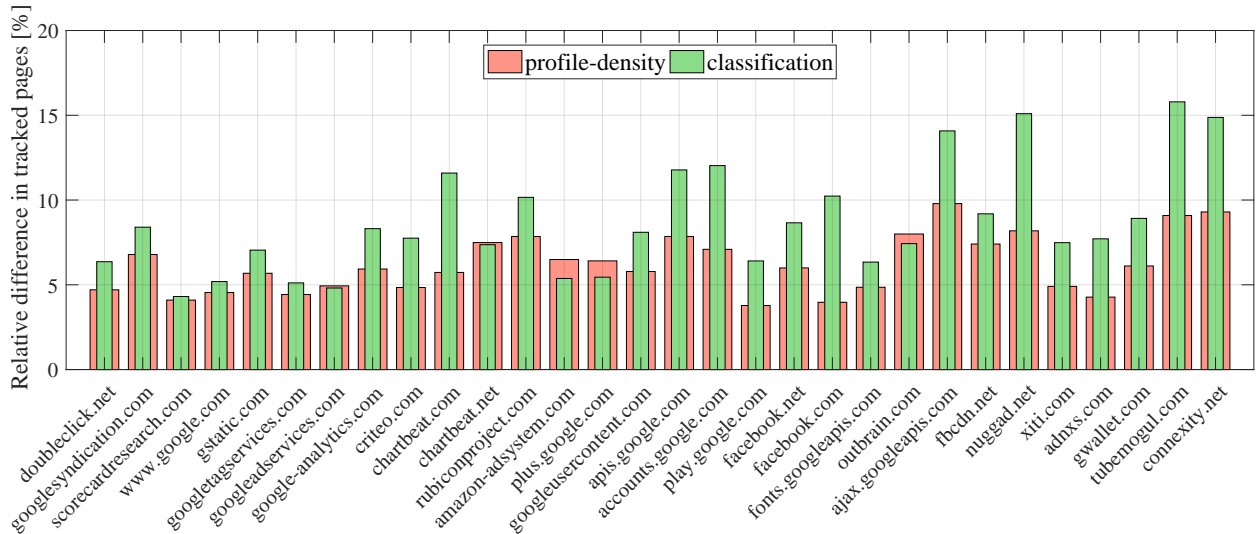


Figure 15: Relative difference in the number of tracked pages per user, between a non-paying subscriber and a non-subscriber.

65.2% and 61.9% of the pages it tracks through third-party network requests, when users respectively choose a profile-density and a classification privacy model. In the case of non-paying subscribers, we notice slightly smaller average values, namely, 59.3% and 59.8% of blocked pages for profile density and classification.

Obviously, this is a limitation of our approach, but the impact of these blocked pages on the Web economy is much less severe than that caused by the current ad-blocking and anti-tracking solutions, which eliminate 100% of any form of tracking and advertising when the maximum privacy level is required. Here, for the same level of privacy, those trackers which do not want to adhere to our fairer, privacy-protecting tracking model are blocked (on average) on up to 65.2% of the pages where they are present. However, our approach is far more flexible than the current solutions¹⁶, and allows configuring multiple points of operation within the optimal trade-off between privacy and money, for which said blocking effect is reduced significantly.

Before proceeding with the scenario B, we would like to point out an important difference between being a non-paying subscriber and a non-subscriber. We conjectured in Sec. 4.3.4 that a subscriber would be in a better position to estimate a user’s activity, and thus decide whether to bid or not for them. Fig. 15 plots the difference in observed activity, averaged over all users, between a non-paying subscriber and a non-subscriber. We observe an average relative difference of 14.9% and 16.1% for the profile-density and the classification models, and confirm that, at least in our data set, *all* trackers would benefit from a greater number of tracked pages per user.

Having examined the case where trackers do not compensate users for tracking, next we turn to the scenario B where they stand ready to pay whatever is necessary to learn users’ actual browsing profiles.

Fig. 16 shows the relative increase in tracked pages —again averaged over all users— for this latter case, which unsurprisingly yields much better results; by “better” we mean that trackers find more opportunities, compared to scenario A, to follow users’ visits throughout the Web. Evidently, when $\mu = \mu_{\max}$, ad platforms and trackers are able to fully track them, which means that their observed clickstreams coincide with the actual clickstream. This implies, on the one hand, that trackers will not be blocked on any of the pages where they are present, that is, $\beta_{\text{crit}} = 0$ for all users. And on the other hand, our approach will allow them to serve ads that can be personalized to accurate browsing interests. A quick look to the data reveals the differences between both scenarios: the relative increase in tracked pages, averaged over the 30 trackers shown, is 16832% in this new case B. We also observe that the maximum value is attained for the tracker domain **nr-data.net** and yields 45393%, which doubles the figure shown for the scenario A.

Continuing with our analysis of the enhanced tracking capabilities, now we explore the content categories and therefore the advertisers that would benefit the most from the proposed pay-per-tracking model. Fig. 17

¹⁶Recall that ad blockers and anti-trackers only provide users with an on/off privacy configuration.




Figure 16: Relative increase in the number of tracked pages in the scenario B.

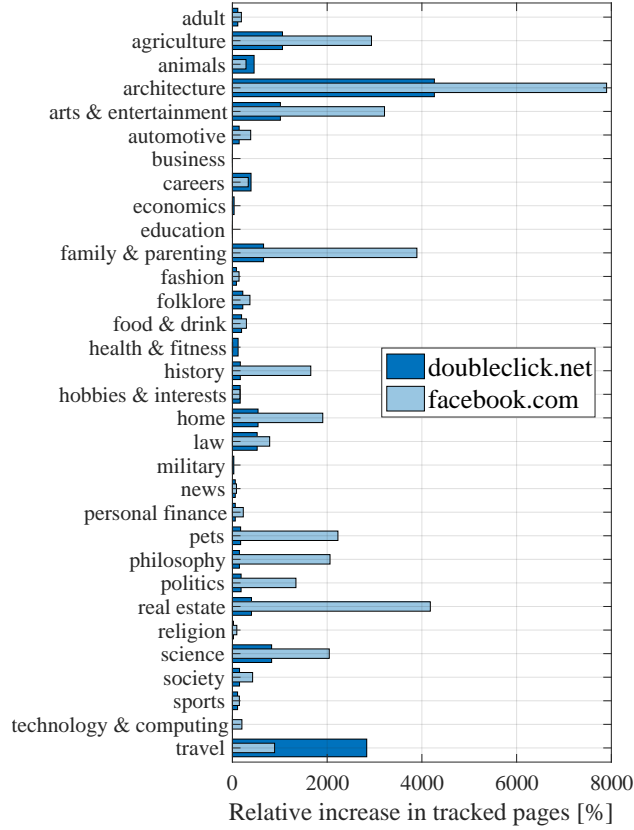


Figure 17: Relative increase in tracked pages per interest category.

shows again the relative increase in tracked pages for case B, but broken down by advertiser type. This figure only shows the results corresponding to **doubleclick.net** and **facebook.com**, which are the most prevalent tracker domains in our data set. Among other aspects, it is worth noting that **doubleclick.net** would be able to increase its presence on travel pages by 2803%, thus improving the effectiveness of personalized travel-related ads. On the other hand, by having access to the actual clickstream, **facebook.com** would




Figure 18: Total number of ads delivered and blocked per tracker domain. Trackers are sorted in decreasing order of served ads. We only display the top 30 ad platforms.

boost their tracking capabilities in “agriculture”, “family & parenting” and “real estate” pages by 7900, 3895 and 4178%, respectively. Those categories less affected by the proposed tracking service turn out to be “business”, “economics”, “education” and “health & fitness”.

Next, we analyze the number of ads that non-paying subscribers could have lost in the scenario A, as a result of the mimicking effect described in Sec. 4.3.4. We estimate the number of ads served by an ad platform on the basis of the iframes contained in the pages visited by users. In total, we found 835 different domains delivering iframes to our 82 users. Since the presence of an iframe may not always imply the delivery of an ad, we used the list of advertising domains from Mozilla Focus project, and their partner Disconnect [24], to select those iframes exclusively serving ads. After this filtering, the number of domains and ads became 110 and 45 710, respectively.

In Fig. 18 we show the number of ads of the 30 domains with more delivered ads. Specifically, we represent the number of ads that would be displayed in the current advertising model, as well as the number of those ads that would be blocked in the scenario A, for a classification privacy model. The figure shows the expected connection between the number of tracked pages and the number of displayed ads; ad platforms are sorted in decreasing order of total tracked pages in Fig. 16. We find that **doubleclick.net**, **facebook.com**, **googlesyndication.com** and **yimg.com** are the domains serving more ads. Remarkably enough, the ads delivered by these four ad platforms account for the 72.2% of all ads in our data set.

The results shown in this figure are clearly consistent with the data provided in Fig. 14 about the percentage of blocked pages in this scenario. For example, for the tracker domains **doubleclick.net** and **facebook.com**, we observe 11 438 and 9 554 blocked ads, which roughly approximates to the 90% and 86.7% of blocked pages seen in that figure.

Our last experiments of this subsection aim at exploring the loss in *ad personalization* a tracker experiences in the scenario A, i.e., when we move from the current tracking model to the new one, and the (subscribed) tracker does not want to pay users for tracking. In this situation, the question that we face is, to what extent the initial profiles degrade profile-based ad targeting.




Figure 19: Average relative cosine variation per tracker domain.

The remainder of this section examines the similarity of the tracked profiles with respect to the real distribution in this scenario A. We use the cosine distance, defined as

$$d(p, q) = 1 - \frac{\langle p, q \rangle}{\|p\| \|q\|},$$

to quantify this (dis)similarity between profiles, and accordingly define the *relative cosine variation* as

$$\rho_v = \frac{d(p, q) - d(q', q)}{d(q, q')}.$$

Clearly, large values of ρ_v reflect a situation in which interest-based ad personalization is less effective than in the current advertising model. On the other hand, small or negative values of this quantity suggest similar or improved profile ad targeting.

Fig. 19 shows, for each tracker, the relative cosine variation averaged over all users. We observe a cosine variation of 77.3% over the 733 trackers included in our dataset, and enormous values of this ρ_v for **doubleclick.net** (2831%) and **scorecardresearch.com** (2321%). This latter observation seems to be consistent with the fact that those are two of the most prevalent trackers in our analysis, which suggests that their profiles q' are much closer to the actual ones than to the initial distributions. On the other hand, although not shown in this figure, we evidenced negative values of ρ_v for 48 tracker domains. Not entirely unexpected, these results were observed for those trackers with a smaller presence on the Web. The minimum observed value was -17.28% .

In short, the main conclusion that can be drawn from Fig. 19 is that, in general, the scenario A would make trackers (non-paying subscribers) observe profiles that are more distant from the real ones, when compared to the current advertising model; obviously, this translates into worse personalization performance. Notwithstanding, those trackers with very limited tracking capabilities might even enhance the effectiveness of their ad-targeting algorithms.

5.4.2. Scenario C

This subsection delves into the scenario C, where we assume our pay-per-tracking model is adopted and all trackers are adhered to it. Under this assumption, we compare the case when users are not rewarded for being tracked, to the case when they are offered $\mu \in (0, \mu_{\max}]$ for it.

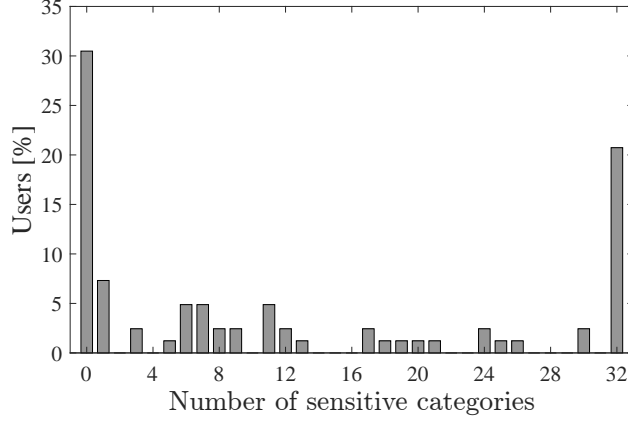


Figure 20: Probability distribution of the number of sensitive categories declared by users.




Figure 21: Percentage of users declaring a specific category as sensitive.

Unlike the previous scenarios A and B, here we need to know the weights $(w_i)_{i=1}^n$, that is, the amount of money each user would ask for completely disclosing their browsing interests in each category. In our experimental analysis, we estimate such weights through the categories they declared as sensitive. In particular, we suppose that all them assign a same weight w_s to their sensitive categories, and a same weight w_{ns} to their non-sensitive ones. In these experiments we assume $w_s = 3w_{ns}$ for all users. On the other hand, for convenience we shall use a normalized reward $\bar{\mu} = \frac{\mu}{\mu_{\max}}$.

Fig. 20 portrays the number of categories deemed as sensitive by our 82 users. Quite surprisingly, we note that 30.5% of them regarded all categories as non-sensitive, while 20.7% of users considered the opposite. This yields an average of 12.2 sensitive categories per user, and gives an idea of the potential cost to a tracker of unveiling a user’s actual profile. Fig. 21 shows the specific categories that have been declared as sensitive, and for which trackers will need to pay more. As expected, “adult”, “religion” and “health & fitness” are at the top of this list.

We start our analysis by examining the extent to which our collaborative-masking scheme contributes to privacy protection. For this purpose, first, we explore how a particular user in our data set benefits from the application of an optimal profile disclosure strategy. And secondly, we consider the whole population




Figure 22: Apparent profiles for different values of $\bar{\mu}$ and for the privacy classification model.

of users and assess the relative reduction in privacy risk when these users are offered the same normalized reward.

To conduct our first experiments, we chose a particular user from our data set. The user considered in these experiments browsed 1 133 pages and selected the categories “adult”, “religion” and “health & fitness” as sensitive. In Fig. 22 we represent the apparent profile of this user for different values of $\bar{\mu}$, in the case of a privacy classification model. When $\bar{\mu} = 0$, any tracker observes an apparent profile t that coincides with the profile representative of cluster number 9, which implies, by virtue of the chosen privacy model, that $\mathcal{R} = 0$. When offered one third of μ_{\max} , things change but fortunately not too much. Our masking algorithm recommends that the user disclose almost their complete interest value in the categories 9 (“economics”), 25 (“politics”) and 27 (“religion”), and half of it in the category 20 (“military”). The disclosure of these categories approximately amounts to the offered money. But remarkably enough, the attained privacy risk *just* represents 0.16% of the final privacy risk $D(q||p)$.

This riveting effect is also observed for $\bar{\mu} \simeq 0.6667$. Our system fully reveals the values of the categories 3, 11, 14, 17, 21 and 28, but this leads to a reduction of only 7.8% of the final privacy risk. Finally, when $\bar{\mu} = 1$ no perturbation takes place and the apparent profile t represented in Fig. 22(d) actually corresponds to the genuine user profile q . All this information is captured in Fig. 23, where we plot the privacy-money function (1), that is, the function modeling the optimal trade-off between privacy and economic reward.

The second set of experiments contemplates a scenario where all users are offered a same $\bar{\mu}$. Note that, in practice, each user would very likely be given a different compensation. Under this assumption, Fig. 24




Figure 23: Optimal trade-off between privacy and money for a particular user in our dataset.

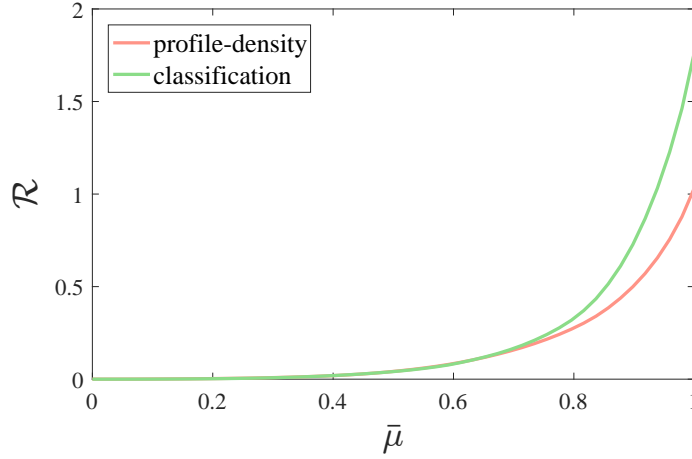


Figure 24: Privacy risk averaged over all users in the scenario C.

shows the average privacy protection achieved by those users, for the classification and the profile-density models. These results were obtained by computing numerically¹⁷ the solution to the problem (1).

Consistently with the results shown for a single user, we observe a relatively small increase of \mathcal{R} for values of $\bar{\mu}$ lower than 0.6. For that particular value, for example, users achieve a privacy risk which is only 4.4% and 7.5% of the final privacy risk (i.e., when they disclose their real profiles) for the classification and the profile-density models, respectively. As commented above, this effect has an immediate and interesting consequence: users may benefit from being tracked *without* significantly harming their privacy. This is true for small to moderate gains. However, for large $\bar{\mu}$, an increase in the economic reward by $\Delta\bar{\mu}$ translates into a increase in privacy risk $\Delta\mathcal{R} > \Delta\bar{\mu}$. That is, users see how their privacy is severely compromised for relatively large compensations. Lastly, we observe that classification yields worse results in terms of privacy risk than the profile-density model. This may be justified by the fact that the initial profiles defined in the latter model are closer to the actual profiles.

Having examined user privacy, next we focus on the benefits that our approach may bring to those ad companies which decide to pay users. In particular, we explore the benefits in terms of tracking and profile-based targeting. Before proceeding, we would like to stress that the results provided next apply to any tracker in our dataset.

¹⁷The numerical method chosen is the interior-point optimization algorithm [32] implemented by the Matlab R2016a function `fmincon`.




Figure 25: Relative increase in tracked pages for any tracker in the scenario C.




Figure 26: Average cosine distance between an actual browsing profile q and the perturbed profile t for some economic reward $\bar{\mu}$.

In Fig. 25, we show the relative increase in tracked pages averaged over all users, when trackers go from observing p to t for some positive $\bar{\mu}$. As with scenarios A and B, we see enormous rates of increase. However, these rates are only observed here for relatively large values of $\bar{\mu}$. For instance, in order for a tracker to increase the number of tracked pages by 10%, it will need to pay $\bar{\mu} \simeq 0.49$ on average per user. On the other hand, for the same reasons provided in Fig. 24, the profile-density model yields worse results in terms of tracked pages when $\bar{\mu}$ approaches 1. Clearly, since this model assumes initial profiles that are closer to the original ones, we may expect lower values of β_{crit} . In short, in our dataset a tracker will require offering large compensations to significantly increase its tracking capabilities.

In our last figure, Fig. 26, we plot the average cosine distance between a user's apparent profile and their actual distribution, $d(t, q)$, again for different values of $\bar{\mu}$. Recall that the cosine distance ranges from 0, meaning exactly the same, to 1, meaning orthogonality. As expected, larger compensations elicit smaller (average) cosine distances between the profiles t and q of each user in our dataset. Analogously to our previous observations for privacy risk and tracking enhancement, the figure also shows a slow decay of our measure of profile accuracy with $\bar{\mu}$. Clearly, this indicates that trackers will have to pay amounts closer to μ_{max} for them to serve completely effective profile-based ads. From the standpoint of a nonpaying tracker, a positive aspect is that the profile-density model provides initial profiles which are not very dissimilar to the actual ones. Evidently, because the initial profiles chosen in the classification model are farther away from the corresponding real distributions, we see larger values of average cosine distance for this model.

In closing, for small relative compensations (e.g., $\bar{\mu} < 0.2$), user privacy remains almost unaltered and trackers get between 0 and 7 percent more of pages to track. However, this slightly enhanced tracking cannot be used to deliver profile-based ads, since the observed profiles differ between 0.27 to 0.47 in cosine distance for the profile-density and classification models, respectively. On the other extreme, when $\bar{\mu} \rightarrow 1$, users disclose very accurate profiles, trackers increase their opportunities to track by 3 023% (profile-density) and 4 979% (classification), and ads may be targeted to a user’s interests with unprecedented accuracy. In a nutshell and although these experiments obviously dependent on the underlying dataset, the results yielded by this analysis provide convincing evidence of the benefits of the proposed pay-per-tracking service to both users and ad companies.

6. Conclusions and Future Work

Ad blockers and anti-trackers are threatening the business model that has sustained the Web so far. These technologies emerged as a response to the proliferation of invasive ads and to the intrusiveness of tracking and profiling practices [64]. However, blocking all tracking, and by extension all ads, has not improved the situation but made it considerably worse.

In this paper, we propose a new tracking model which places privacy at the forefront of design, shifts the balance of power between users and ad companies, and at the same time, it may safeguard the Internet economic model even though when these companies are not disposed to cooperate. Our pay-per-tracking service may preserve the current advertising model in the sense that non-cooperative ad platforms can still track and deliver ads. Nevertheless, behavioral targeting, the most effective form of advertising, is no longer effective under this scheme.

We achieve this by allowing users to show an initial profile when no reward is offered to them. This profile is chosen thanks to the collaboration among users of our service, which permits them to find out how the other profiles are and, accordingly, decide how they want to be seen in such worst case. To make this decision, our model does not restrict to a particular privacy metric or model. Instead, we allow users to select the most appropriate one for their requirements. Despite this, we consider two privacy models for the purpose of illustration, one based on the commonality of profiles, and the other based on the proximity to a certain group of users.

For a given compensation, the disclosed profile is obtained as the convex combination between the initial profile and the actual one. We formulate the problem of disclosing a user’s profile as a multiobjective optimization problem that characterizes the optimal trade-off between privacy and money. The result is a perturbation rule that helps users protect their privacy, and allows them to find the optimal exchange of privacy for money.

We propose a system architecture that implements this perturbative rule in real practice, as users navigate the Web, without affecting the browsing experience. The core element of this architecture is a module that attempts to emulate the optimal apparent profile, on the basis of the pages visited by the user, by allowing or blocking tracking on those pages.

We carry out a thorough experimental analysis assessing the extent to which our approach, on the one hand, may contribute to privacy protection and at the same time tracking enhancement, and on the other it may negatively affect the Web economy.

Among other results, our findings show that, in the scenario A, non-paying subscribers expand their tracking capabilities enormously, far beyond their observed clickstream. Yet, the expanded capabilities come at the cost of blocking an average of 59.8% of the pages and ads where they are present. Besides, the resulting profiles exhibit an average relative variation over 77%, which means that the initial profiles are further away from the actual ones than the profiles that could have been built before our pay-per-tracking model is deployed.

For non-subscribers, things get worse in terms of number of blocked pages, averaging up to 65.2% over all trackers. In addition to lower rates of blocked pages, a non-paying subscriber is in a better position to estimate a user’s activity. For subscribers which fully compensate users for tracking, we observe high rates of increase in the number of tracked pages, when compared to the current tracking model. On average, trackers enhance their tracking capabilities by 16 832%.

In the scenario C, for values of $\bar{\mu}$ lower than 0.6, we notice an increase of privacy risk by just 7.5% of the final privacy risk (i.e., when the actual profile is fully disclosed), and an improvement in profile similarity and number of tracked pages approximately by 10% and 55%, respectively. On the other hand, for values of $\bar{\mu}$ close to 1, trackers boost their presence on the Web by 4979% with respect to the case when $\bar{\mu} = 0$, and privacy risk and profile-based personalization reach maximum levels.

Among other aspects, future research should investigate some of the assumptions made in this work. Since we acknowledge that user profiles may vary significantly over time, we need to consider this fact in order to periodically update not only users' profiles but also the profiles available to ad platforms and trackers. Another strand of future work will investigate the case when users may try to re-negotiate tracking with ad platforms, that is, when users partly or fully disclose their profiles for a certain μ , and afterwards they may wish to agree with the ad platform in question on a $\mu' \neq \mu$.

Acknowledgment

This work was partly funded by the European Commission through the H2020 project "CLARUS". J. Parra-Arnau is the recipient of a Juan de la Cierva postdoctoral fellowship, FJCI-2014-19703, from the Spanish Ministry of Economy and Competitiveness.

APPENDIX

Appendix A. Clustering Data

We have applied Lloyd's algorithm [59] to group users into 10 clusters. This appendix shows in Table A.5 the centroid, or representative profile, of each cluster.

References

- [1] Adblock plus, accessed on 2015-10-22.
URL <https://adblockplus.org>
- [2] Citizenme, accessed on 2016-04-18.
URL <https://www.citizenme.com/>
- [3] Clickstream or clickpath analysis, accessed on 2015-03-27.
URL <http://www.opentracker.net/article/clickstream-or-clickpath-analysis>
- [4] COIN-OR Interior Point OPTimizer, accessed on 2015-09-17.
URL <https://projects.coin-or.org/Ipopt>
- [5] Consumer opt-out, Tech. rep., Netw. Advertising Initiative, accessed on 2015-03-19.
URL <http://www.networkadvertising.org/choices>
- [6] Datacoup, accessed on 2016-04-18.
URL <https://datacoup.com/>
- [7] Datawallet, accessed on 2016-04-18.
URL <https://www.datawallet.io/>
- [8] Disconnect, accessed on 2015-11-19.
URL <https://disconnect.me/>
- [9] Ghostery, accessed on 2016-04-08.
URL <https://www.ghostery.com>
- [10] Iab tech lab publisher ad blocking primer, Tech. rep., accessed on 2016-03-13.
URL http://www.iab.com/wp-content/uploads/2016/03/IABTechLab_Publisher_AdBlocking_Primer.pdf
- [11] Interactive advertising bureau, accessed on 2015-09-11.
URL <http://www.iab.com>
- [12] Lightbeam, accessed on 2015-09-24.
URL <https://www.mozilla.org/en-US/lightbeam/>
- [13] The official easylist website, accessed on 2015-10-22.
URL <https://easylist.adblockplus.org>
- [14] Privacy badger, accessed on 2015-09-24.
URL <https://www.eff.org/es/node/73969>
- [15] Cisco service control online advertising solution guide, Tech. rep., Cisco Syst. (Jan. 2009).

Category	Centroid									
	1	2	3	4	5	6	7	8	9	10
1	0.000	0.431	0.000	1.099	1.700	0.005	59.261	0.038	0.011	0.000
2	0.022	0.000	0.000	0.000	0.000	0.008	0.000	0.000	0.016	0.000
3	0.326	0.507	5.043	0.414	3.390	0.474	0.238	0.376	0.658	0.000
4	0.592	1.222	0.596	0.452	0.400	2.242	0.318	1.121	0.483	0.573
5	4.557	22.048	2.391	7.263	39.513	17.645	2.915	5.585	5.221	1.755
6	0.016	0.061	0.310	0.351	0.155	0.189	0.000	0.934	0.000	24.585
7	4.413	3.100	2.050	3.837	8.116	5.080	0.079	2.950	22.675	2.155
8	0.528	0.989	0.146	0.882	0.000	0.237	0.000	0.985	3.713	0.847
9	0.367	0.944	2.925	0.485	0.019	0.905	0.000	1.862	0.159	15.877
10	2.204	1.406	0.144	0.893	0.000	3.391	0.238	3.862	0.614	0.000
11	0.023	0.029	0.000	0.069	0.038	0.062	0.000	0.032	0.005	0.000
12	0.220	0.227	0.144	0.087	0.046	0.177	0.556	0.150	1.524	0.000
13	0.005	0.022	0.000	0.004	0.000	0.105	0.000	0.053	0.098	0.000
14	0.413	0.372	0.000	0.244	0.786	1.271	0.159	1.195	0.479	0.000
15	0.958	1.338	0.439	0.607	0.872	4.798	1.590	1.618	4.630	0.299
16	0.077	0.081	0.000	0.013	0.000	0.028	0.000	0.015	0.006	0.000
17	15.760	2.521	0.828	1.916	0.687	1.278	0.808	3.760	28.308	0.971
18	0.069	0.053	0.110	0.033	0.094	0.030	0.000	0.028	0.000	0.000
19	0.449	0.395	0.144	0.069	0.059	0.284	0.000	0.152	0.531	0.000
20	0.154	0.322	0.000	0.157	0.182	0.568	0.000	0.764	0.288	0.000
21	1.292	0.422	0.072	0.290	0.398	3.597	0.000	0.832	1.849	18.469
22	0.623	1.604	1.497	0.878	0.230	2.780	0.636	2.939	0.402	0.541
23	0.393	0.015	0.000	0.019	0.000	0.049	0.000	0.025	0.068	0.000
24	0.000	0.007	0.000	0.050	0.000	0.099	0.000	0.045	0.000	0.000
25	1.145	0.444	0.914	0.650	0.030	4.527	0.159	1.059	0.169	0.560
26	0.003	0.000	0.000	0.000	0.000	0.020	0.000	0.000	0.000	0.000
27	0.313	0.194	0.310	0.174	0.080	0.365	0.079	0.382	0.193	0.000
28	8.828	10.281	1.777	7.129	28.281	13.049	8.810	11.283	5.200	9.483
29	19.918	1.680	1.401	52.772	5.093	9.731	10.056	7.811	3.415	2.710
30	1.778	1.212	2.239	0.804	1.200	2.976	1.789	2.530	1.273	16.425
31	34.436	47.799	76.520	15.117	8.631	21.522	12.308	47.360	17.878	4.750
32	0.116	0.276	0.000	3.241	0.000	2.508	0.000	0.256	0.134	0.000

Table A.5: Interest values per category, expressed in percentage terms, for the 10 centroids used in our experimental evaluation.

- [16] Topline u.s. web data for march 2010, Tech. rep., accessed on 2015-02-19 (Mar. 2010).
URL <http://www.nielsen.com/us/en/insights/news/2010/nielsen-provides-topline-u-s-web-data-for-march-2010.html>
- [17] Adblock plus user survey results, part 3, Tech. rep., Eyeo, accessed on 2015-07-11 (Dec. 2011).
URL <https://adblockplus.org/blog/adblock-plus-user-survey-results-part-3>
- [18] The state of online advertising, Tech. rep., Adobe, accessed on 2015-09-11 (2012).
URL <http://www.adobe.com/aboutadobe/pressroom/pdfs/Adobe.State.of.Online.Advertising.Study.pdf>
- [19] Firefox interest dashboard, accessed on 2015-05-02 (Nov. 2014).
URL <https://www.mozilla.org/en-US/firefox/interest-dashboard/>
- [20] US programmatic ad spend tops \$10 billion this year, to double by 2016, Tech. rep., eMarketer (Oct. 2014).
URL <http://www.emarketer.com/Article/US-Programmatic-Ad-Spend-Tops-10-Billion-This-Year-Double-by-2016/1011312>
- [21] The cost of ad blocking, Res. rep., PageFair (Aug. 2015).
- [22] State of privacy report 2015, Tech. rep., Symantec, accessed on 2015-05-10 (Feb. 2015).
URL <https://www.symantec.com/content/en/us/about/presskits/b-state-of-privacy-report-2015.pdf>
- [23] Tracking preference expression (DNT), Tech. rep. (Aug. 2015).
URL <http://www.w3.org/TR/tracking-dnt/>
- [24] Efficient & optional filtering of domains in Recursor 4.0.0, <http://blog.powerdns.com/2016/01/19/efficient-optional-filtering-of-domains-in-recursor-4-0-0/>, accessed 2016-02-21 (2016).
- [25] J. P. Achara, J. Parra-Arnau, C. Castelluccia, *MyTrackingChoices*: Pacifying the ad-block war by enforcing user privacy preferences, in: Proc. Annual Workshop Econ. Inform. Secur. (WEIS), 2016, to appear.
- [26] J. P. Achara, J. Parra-Arnau, C. Castelluccia, *Mytrackingchoices* (Feb. 2016).
URL <https://chrome.google.com/webstore/detail/mytrackingchoices/fmonkjimgifgcgeodhhgbfoncmjclka>
- [27] G. Aisch, W. Andrews, J. Keller, The cost of mobile ads on 50 news websites, accessed on 2015-12-19 (Oct. 2015).
URL <http://www.nytimes.com/interactive/2015/10/01/business/cost-of-mobile-ads.html>
- [28] M. Aly, A. Hatch, V. Josifovski, V. K. Narayanan, Web-scale user modeling for targeting, in: Proc. Int. WWW Conf., ACM, 2012, pp. 3–12.
- [29] L. C. Andrei, *Money and Market in the Economy of All Times: another world history of money and pre-money based economies*, Xlibris Publishing, 2011.
- [30] Audience Buying Guide 2011, accessed on 2015-03-27.
URL <http://brandedcontent.adage.com/audiencebuyingguide2011/network.php?id=12>
- [31] H. Beales, The value of behavioral targeting, Tech. rep., Netw. Advertising Initiative, accessed on 2016-01-15 (Mar. 2010).
URL http://www.networkadvertising.org/pdfs/Beales_NAI_Study.pdf
- [32] S. Boyd, L. Vandenberghe, *Convex Optimization*, Cambridge University Press, Cambridge, UK, 2004.
- [33] R. Cookson, Google, Microsoft and Amazon pay to get around ad blocking tool, accessed on 2014-03-10 (Feb. 2015).
URL <http://www.ft.com/cms/s/0/80a8ce54-a61d-11e4-9bd3-00144feab7de.html>
- [34] T. M. Cover, J. A. Thomas, *Elements of Information Theory*, 2nd ed., Wiley, New York, 2006.
- [35] J. Davidson, Forget facebook, meet the company that will pay you for your personal data, accessed on 2016-04-12 (Sep. 2014).
URL <http://time.com/money/3001361/datacoup-facebook-personal-data-privacy/>
- [36] J. Domingo-Ferrer, J. Soria-Comas, O. Ciobotaru, Co-utility: self-enforcing protocols without coordination mechanisms, in: Proc. Int. Conf. Ind. Eng., Oper. Manage. (IEOM), IEEE Comput. Soc., 2015, pp. 1–17.
- [37] Y. Elovici, C. Glezer, B. Shapira, Enhancing customer privacy while searching for products and services on the World Wide Web, *Internet Res.* 15 (4) (2005) 378–399.
- [38] Y. Elovici, B. Shapira, A. Maschiach, chap. A New Privacy Model for Web Surfing, Springer-Verlag, 2002, pp. 45–57.
- [39] Y. Elovici, B. Shapira, A. Maschiach, A new privacy model for hiding group interests while accessing the Web, in: Proc. Workshop Priv. Electron. Soc., ACM, Washington, DC, 2002, pp. 63–70.
- [40] Y. Elovici, B. Shapira, A. Meshiach, Cluster-analysis attack against a private Web solution (PRAW), *Online Inform. Rev.* 30 (2006) 624–643.
- [41] S. Englehardt, The hidden perils of cookie syncing, accessed on 2014-12-10 (Aug. 2014).
URL <https://freedom-to-tinker.com/blog/englehardt/the-hidden-perils-of-cookie-syncing/>
- [42] A. Erola, J. Castellà-Roca, A. Viejo, J. M. Mateo-Sanz, Exploiting social networks to provide privacy in personalized Web search, *J. Syst., Softw.* 84 (10) (2011) 1734–745.
URL <http://www.sciencedirect.com/science/article/pii/S0164121211001117>
- [43] U. Feige, A. Fiat, A. Shamir, Zero-knowledge proofs of identity, *J. Cryptology* 1 (2) (1988) 77–84.
- [44] S. Gauch, M. Speretta, A. Chandramouli, A. Micarelli, chap. User profiles for personalized information access, Springer-Verlag, 2007, pp. 54–89.
- [45] J. Goldberger, S. Gordon, H. Greenspan, An efficient image similarity measure based on approximations of kl-divergence between two gaussian mixtures, in: Proc. Int. Conf. Comput. Vision (ICCV), 2003, pp. 487–493.
- [46] A. Ha, Datawallet pays users to share their online data with businesses, accessed on 2016-04-10 (Jul. 2015).
URL <http://techcrunch.com/2015/07/10/datawallet-seed-funding/>
- [47] M. Halkidi, I. Koutsopoulos, A game theoretic framework for data privacy preservation in recommender systems, in: Proc. European Mach. Learn., Prin., Pract. Knowl. Disc. Databases (ECML PKDD), Springer-Verlag, 2011, pp. 629–644.
- [48] A. Hundepool, J. Domingo-Ferrer, L. Franconi, S. Giessing, E. S. Nordholt, K. Spicer, P.-P. de Wolf, *Statistical Disclosure Control*, Wiley, 2012.

- [49] Q. Huo, W. Li, A dtw-based dissimilarity measure for left-to-right hidden markov models and its application to word confusability analysis, in: Proc. Interspeech, 2006, pp. 2338–2341.
- [50] R. Joe, The book of fraud: A marketer's guide to bots, fake domains and other dirty deeds in online advertising (Dec. 2014).
URL <http://adexchanger.com/online-advertising/the-book-of-fraud>
- [51] S. G. Johnson, NLOpt nonlinear-optimization package, accessed on 2015-09-16.
URL <http://ab-initio.mit.edu/nlopt>
- [52] A. Kae, K. Kan, V. K. Narayanan, D. Yankov, Categorization of display ads using image and landing page features, in: Proc. ICDM Workshop Large-Scale Data Min.: Theory, Appl., ACM, 2011, pp. 1–8.
URL <http://doi.acm.org/10.1145/2002945.2002946>
- [53] T. Kawaja, Display LUMAscape, accessed on 2015-09-23.
URL <http://www.lumapartners.com/lumascape/display-ad-tech-lumascape>
- [54] T. Kuflik, B. Shapira, Y. Elovici, A. Maschiach, Privacy preservation improvement by learning optimal profile generation rate, in: User Modeling, vol. 2702 of Lecture Notes Comput. Sci. (LNCS), Springer-Verlag, 2003, pp. 168–177.
- [55] P. Laperdrix, W. Rudametkin, B. Baudry, Mitigating browser fingerprint tracking: multi-level reconfiguration and diversification, in: Proc. Int. Symp. Softw. Eng. Adap. Self-Manag. Syst. (SEAMS), 2015.
- [56] P. Laperdrix, W. Rudametkin, B. Baudry, Beauty and the beast: Diverting modern web browsers to build unique browser fingerprints, in: Proc. IEEE Symp. Secur., Priv. (SP), IEEE Comput. Soc., 2016.
- [57] G. Lindsay, B. Woods, J. Corman, Smart homes and the internet of things, Tech. rep., Atlantic Council, accessed on 2016-04-12 (Mar. 2016).
URL <http://www.atlanticcouncil.org/publications/reports/smart-homes-and-the-internet-of-things>
- [58] B. Liu, A. Sheth, U. Weinsberg, J. Chandrashekar, R. Govindan, Adreveal: Improving transparency into online targeted advertising, in: Proc. Hot Topics in Netw., ACM, 2013, pp. 12:1–12:7.
- [59] S. P. Lloyd, Least squares quantization in PCM, IEEE Trans. Inform. Theory IT-28 (1982) 129–137.
- [60] D. Lyon (ed.), Surveillance as Social Sorting: Privacy, Risk and Automated Discrimination, Routledge, 2002.
- [61] B. Markines, C. Cattuto, F. Menczer, D. Benz, A. Hotho, G. Stum, Evaluating similarity measures for emergent semantics of social tagging, in: Proc. Int. WWW Conf., ACM, 2009, pp. 641–650.
- [62] G. Marvin, Consumers now notice retargeted ads, Tech. rep., Marketing Land, accessed on 2015-08-12 (Dec. 2013).
URL <http://marketingland.com/3-out-4-consumers-notice-retargeted-ads-67813>
- [63] Q. Mei, C. Zhai, Discovering evolutionary theme patterns from text: an exploration of temporal text mining, in: Proc. ACM SIGKDD Int. Conf. Knowl. Disc., Data Min. (KDD), ACM, 2005, pp. 198–207.
- [64] W. Melicher, M. Sharif, J. Tan, L. Bauer, M. Christodorescu, P. G. Leon, (Do Not) track me sometimes: Users contextual preferences for web tracking, in: Proc. Int. Symp. Priv. Enhanc. Technol. (PETS), Lecture Notes Comput. Sci. (LNCS), Springer-Verlag, 2016, pp. 1–20.
- [65] B. Morrissey, Forbes starts blocking ad-block users, accessed on 2016-02-29 (Dec. 2015).
URL <http://digiday.com/publishers/forbes-ad-blocking/>
- [66] K. Murphy, The ad blocking wars, accessed on 2015-02-22 (Feb. 2016).
URL <http://www.nytimes.com/2016/02/21/opinion/sunday/the-ad-blocking-wars.html>
- [67] C. Newton, Facebook considers letting users add a tip jar to make money from posts, accessed on 2016-04-25 (Apr. 2016).
URL <http://www.theverge.com/2016/4/19/11455840/facebook-tip-jar-partner-program-monetization>
- [68] L. Olejnik, Measuring the privacy risks and value of web tracking, Ph.D. thesis, Nat. Inst. Res. Comput. Sci., Contr. (INRIA) (Jan. 2015).
- [69] L. Olejnik, T. Minh-Dung, C. Castelluccia, Selling off privacy at auction, in: Proc. IEEE Symp. Netw. Distrib. Syst. Secur. (SNDSS), Internet. Soc., 2014.
- [70] P. A. Olsen, S. Dharanipragada, An efficient integrated gender detection scheme and time mediated averaging of gender dependent acoustic models, in: Proc. European Conf. Speech Commun., Technol. (EUROSPEECH), 2003, pp. 2509–2512.
- [71] D. Olszewski, Fraud detection in telecommunications using kullback-leibler divergence and latent dirichlet allocation, in: Proc. Adap., Nat. Comput. Alg., Springer-Verlag, 2011, pp. 71–80.
- [72] S. Pandey, M. Aly, A. Bagherjeiran, A. Hatch, P. Ciccio, A. Ratnaparkhi, M. Zinkevich, Learning to target: What works for behavioral targeting, in: Proc. Int. Conf. Inform., Knowl. Manage. (CIKM), ACM, 2011, pp. 1805–1814.
- [73] J. Parra-Arnau, Privacy protection of user profiles in personalized information systems, Ph.D. thesis, Tech. Univ. Catalonia (UPC) (Dec. 2013).
URL <https://sites.google.com/site/javierparraarnau/publications/JParra-Arnau-PhDThesis.pdf>
- [74] J. Parra-Arnau, J. P. Achara, C. Castelluccia, MyAdChoices: Bringing transparency and control to online advertising, ACM Trans. WebArXiv preprint.
URL <http://arxiv.org/abs/1602.02046>
- [75] J. Parra-Arnau, D. Rebollo-Monedero, J. Forné, Measuring the privacy of user profiles in personalized information systems, Future Gen. Comput. Syst. (FGCS), Special Issue Data, Knowl. Eng. 33 (2014) 53–63.
URL <http://dx.doi.org/10.1016/j.future.2013.01.001>
- [76] H. Printz, P. Olsen, Theory and practice of acoustic confusability, Comput. Speech, Lang. 16 (1) (2002) 131–164.
- [77] S. Puglisi, D. Rebollo-Monedero, J. Forné, You never surf alone. ubiquitous tracking of users browsing habits, in: Proc. Int. Workshop Data Priv. Manage. (DPM), Lecture Notes Comput. Sci. (LNCS), Vienna, Austria, 2015.
- [78] D. Rebollo-Monedero, J. Forné, Optimal query forgery for private information retrieval, IEEE Trans. Inform. Theory 56 (9) (2010) 4631–4642.

- [79] D. Rebollo-Monedero, J. Parra-Arnau, C. Diaz, J. Forné, On the measurement of privacy as an attacker's estimation error, *Int. J. Inform. Secur.* 12 (2) (2012) 129–149.
URL <http://link.springer.com/article/10.1007/s10207-012-0182-5>
- [80] D. Rogers, How business can gain consumers' trust around data, accessed on 2015-11-03 (Nov. 2015).
URL <http://www.forbes.com/sites/davidrogers/2015/11/02/how-business-can-gain-consumers-trust-around-data/>
- [81] P. Samarati, Protecting respondents' identities in microdata release, *IEEE Trans. Knowl. Data Eng.* 13 (6) (2001) 1010–1027.
- [82] P. Sayer, Adblock extension begins whitelisting “acceptable ads” (Oct. 2015).
URL <http://www.pcworld.com/article/2988838>
- [83] M. J. Schervish, *Theory of Statistics*, Springer-Verlag, New York, 1995.
- [84] B. Shapira, Y. Elovici, A. Meshiach, T. Kuflik, PRAW – The model for P_{Priv}Ate Web, *J. Amer. Soc. Inform. Sci., Technol.* 56 (2) (2005) 159–172.
- [85] R. Shokri, G. Theodorakopoulos, J. Y. L. Boudec, J. P. Hubaux, Quantifying location privacy, in: *Proc. IEEE Symp. Secur., Priv. (SP)*, IEEE Comput. Soc., Washington, DC, USA, 2011, pp. 247–262.
- [86] M. Smith, *Targeted: How Technology Is Revolutionizing Advertising and the Way Companies Reach Consumers*, 1st ed., AMACOM, New York, 2014.
- [87] J. Soria-Comas, J. Domingo-Ferrer, Co-utile collaborative anonymization of microdata, in: *Proc. Int. Conf. Model. Decisions Artif. Intell.*, 2015, pp. 192–206.
- [88] L. Sweeney, *k*-Anonymity: A model for protecting privacy, *Int. J. Uncertain., Fuzz., Knowl.-Based Syst.* 10 (5) (2002) 557–570.
- [89] V. Toubiana, *SquiggleSR* (2007).
URL www.squigglesr.com
- [90] V. Toubiana, A. Narayanan, D. Boneh, H. Nissenbaum, S. Barocas, Adnostic: Privacy preserving targeted advertising, in: *Proc. IEEE Symp. Netw. Distrib. Syst. Secur. (SNDSS)*, 2010, pp. 1–21.
- [91] M. M. Tsang, S. C. Ho, T. P. Liang, Consumer attitudes toward mobile advertising: An empirical study, *Int. J. Electron. Commer.* 8 (3) (2004) 65–78.
- [92] A. Wächter, L. T. Biegler, On the implementation of a primal-dual interior point filter line search algorithm for large-scale nonlinear programming, *Math. Program.* 106 (1) (2006) 25–57.
- [93] J. Yan, N. Liu, G. Wang, W. Zhang, Y. Jiang, Z. Chen, How much can behavioral targeting help online advertising?, in: *Proc. Int. WWW Conf.*, ACM, 2009, pp. 261–270.
- [94] S. Ye, F. Wu, R. Pandey, H. Chen, Noise injection for search privacy protection, in: *Proc. Int. Conf. Comput. Sci., Eng.*, IEEE Comput. Soc., 2009, pp. 1–8.
- [95] YourOnlineChoices.
URL <http://www.youronlinechoices.com/>
- [96] S. Yuan, A. Z. Abidin, M. Sloan, J. Wang, Internet advertising: An interplay among advertisers, online publishers, ad exchanges and web users, arXiv: 1206.1754ArXiv preprint.
- [97] C. Zhu, R. H. Byrd, J. Nocedal, L-bfgs-b: Algorithm 778: L-bfgs-b fortran routines for large scale bound constrained optimization, *ACM Trans. Math. Softw.* 23 (4) (2007) 550–560.