

My Mouse, My Rules

Privacy Issues of Behavioral User Profiling via Mouse Tracking

Luis A. Leiva
University of Luxembourg
Luxembourg
name.surname@uni.lu

Ioannis Arapakis
Telefonica Research
Spain
ioannis.arapakis@telefonica.com

Costas Iordanou
Cyprus University of Technology
Cyprus
costas.iordanou@eecei.cut.ac.cy

ABSTRACT

This paper aims to stir debate about a disconcerting privacy issue on web browsing that could easily emerge because of unethical practices and uncontrolled use of technology. We demonstrate how straightforward is to capture behavioral data about the users at scale, by unobtrusively tracking their mouse cursor movements, and predict user's demographics information with reasonable accuracy using five lines of code. Based on our results, we propose an adversarial method to mitigate user profiling techniques that make use of mouse cursor tracking, such as the recurrent neural net we analyze in this paper. We also release our data and a web browser extension that implements our adversarial method, so that others can benefit from this work in practice.

CCS CONCEPTS

• **Security and privacy** → **Human and societal aspects of security and privacy**; • **Human-centered computing** → Human computer interaction (HCI).

KEYWORDS

Mouse Cursor Tracking; User Profiling; Privacy; Ethics

ACM Reference Format:

Luis A. Leiva, Ioannis Arapakis, and Costas Iordanou. 2021. My Mouse, My Rules: Privacy Issues of Behavioral User Profiling via Mouse Tracking. In *Proceedings of the 2021 ACM SIGIR Conference on Human Information Interaction and Retrieval (CHIIR '21), March 14–19, 2021, Canberra, ACT, Australia*. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3406522.3446011>

1 INTRODUCTION

In the modern Web, privacy is becoming a rare commodity. The recent proliferation of intrusive and privacy-invasive ads has raised serious concerns among users and industry regulatory bodies, with initial user reaction reflected on the swift adoption of ad blocking solutions. In fact, most of the popular browser extensions for Mozilla

Firefox are related to ad blocking and user privacy.¹ While extensions like these have been successful in mitigating the user's exposure to web tracking, they eventually hurt web revenue streams, leading to the so called "tragedy of the commons" [54], where the common resource (user attention) will be depleted due to ad blocking. Luckily, some effort has been put to regulate the web tracking landscape, like self-initiatives from the ad industry that include recommendations for good practices [18] and transparency tools such as AdChoices,² to help users understand why they receive specific ads. Along the same line, privacy-preserving web browsers³ allow users to have more control of their online privacy and, at the same time, a financial incentive by the ads that they receive while surfing the web. Moreover, in 2018 the European Union set in place the new General Data Protection Regulation (GDPR) [96] and the state of California in United States enforced the Consumer Privacy Act [73]. Other countries are also following the same route, yet currently online advertising is ubiquitous. Also, it remains the dominant monetization model on the Web, with constantly increasing growth rates and revenues. This has promoted advancements in user tracking and profiling technologies that allow to serve more relevant ad content to the user, and at a higher premium, known as targeted ads or Online Behavioral Advertising [14, 21].

Web tracking and user profiling rely on mechanisms to uniquely identify and track the user's online behavior over time, including e.g., geolocation, visited pages, search keywords, and social network activity. All of these in order to better understand the user intentions and interests. However, a less known method to profile the user is by means of mouse cursor tracking. This technology has been used successfully to inform usability tests [10], predict user engagement [5] and intent [32, 66], detect searcher frustration [30], and infer user attention to parts of a web page [6], among other tasks. Unfortunately, because mouse cursor tracking can be performed unobtrusively [58] and at scale [41], it has opened the door to a brand new wave of massive tracking campaigns and companies that hide behind laudable objectives, such as providing fine-grained, in-page analytics (e.g., hovered and clicked items, scroll reach, speed of browsing) to the website owners. Interestingly, by tracking the mouse cursor it is possible to profile the user demographics, namely predicting age [104, 105] and gender [50, 78], a piece of valuable personal data that most users are unaware of [20]. With this paper, we want to raise awareness about this fact and reflect on the trade-offs between privacy and technological innovation.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHIIR '21, March 14–19, 2021, Canberra, ACT, Australia

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8055-3/21/03...\$15.00

<https://doi.org/10.1145/3406522.3446011>

¹<https://addons.mozilla.org/en-US/firefox/search/?promoted=recommended&sort=users&type=extension>

²<http://youradchoices.com/>

³See <https://brave.com/> and <https://cliqz.com/>

We investigate and highlight privacy issues that may emerge because of unethical practices and unregulated use of mouse cursor tracking technology. Our contributions begin with an extensive survey of related work on privacy and security while online, and continue with research that performed user profiling via mouse cursor data. We then show how straightforward it is to capture behavioral data and predict demographics information with reasonable accuracy using a few lines of code. Based on our results, we present an adversarial method to mitigate user profiling techniques that make use of mouse cursor tracking, such as the recurrent neural net we propose in Section 3.6.2.

2 SURVEY OF RELATED WORK

To what extent does our online activity reveal who we are? Existing literature related to online privacy provides insights around topics such as information leakage while surfing the web using desktop computers [60, 82, 91, 99] or mobile devices [61, 83]. Other studies report on the entities that collect tracking data [12, 92] and how tracking data are being collected [2, 76]. Overall, this large body of work demonstrates that the digital footprints left by individuals, as they browse websites, may help derive with alarming accuracy personally identifiable information like gender, age, location, or even political orientation. Recent work by White et al. [100] and Gajos et al. [31] could detect neurodegenerative disorders from mouse cursor movements, showing how our “digital phenotypes” could be used as adjunctive screening tools. In this section, we report current evidence on privacy risks in the online setting, as well as predicting demographic attributes from online digital traces.

2.1 Privacy Issues in Web Browsing

There are many ways for tracking the online activity of web users, for example by monitoring the IP addresses or using fingerprinting techniques [1, 45, 70]. However, the cookie-based approach remains to this day the dominant one, since it is fully supported by all web browsers. With a cookie identifying the user’s browser, a third party domain can track the user activity across websites using redirection techniques or providing a free service that makes cross-domain tracking possible, such as the Facebook ‘Like’ button or social media sharing plugins for WordPress. Then, by analyzing the content of the browsed websites, the tracking domain can effortlessly – and at an unprecedented scale – derive users’ intentions and interests, alongside with other sensitive profiling information [43, 84].

There is a plethora of work that capitalize on user profiling based on how we browse. For example, our browsing history is used to detect targeted ads [21, 77] or even identify which attribute or user action triggered a specific ad [55, 101]. Olejnik et al. [75] noticed that, with just 4 visited websites, it is possible to uniquely identify users in 97% of cases. At the same time, Web users are concerned about third-party tracking [68], especially about location access and inferring demographics [102]. Researchers have found that people are likely to take actions to protect their privacy [65], including the payment of a premium fee if needed [94]. And while some websites and tracking companies inform users about their data practices through privacy policies and sometimes provide opt-outs [85], these measures are insufficient.

Today, advancements in web advertising provide new opportunities to trackers and advertisers to extend their visibility. Online ads are rendered dynamically during the load time of the browsed website, most of the time as a result of additional (tracking) JavaScript code that is injected on the fly [44]. This opens the door to more sophisticated tracking techniques, among which we find mouse cursor tracking to be an underestimated one. In the following section, we highlight how this technology has been used in various scenarios and how pervasive it has become. In fact, today most websites include analytics scripts, and a large number of them contains some mouse tracking script [29].

Users consider ad targeting useful because it highlights relevant information, but at the same time they find the underlying data collection alarming [97] and invasive [67]. Plane et al. [79] found that users were more concerned if an ad was targeted based on demographics, such as age, gender, or race, than based on interests. Overall, users do not want targeted advertising when they are made aware of the data collection methods employed by the advertisers [95], and consider targeting based on demographics to be discriminatory [93]. We, therefore, hypothesize that it could be possible to derive demographics information from mouse movements at scale, and that a privacy issue may emerge if people follow unethical practices and make an uncontrolled use of the technology.

2.2 Mouse Cursor Tracking

What can a mouse cursor tell us more? Almost 20 years ago Chen et al. [22] raised this question and found a relationship between gaze position and cursor position during web browsing. Mueller and Lockerd [71] investigated the use of mouse tracking to visualize and (manually) infer the users’ interests. Since then, researchers have noted the utility of mouse cursor analysis as a low-cost and scalable proxy of eye gaze [39, 72]. Several works have investigated closely the utility of mouse cursor data in web search [6, 24, 63] and web page usability evaluation [9, 10, 56], two of the most prominent use cases of this technology. Mouse biometrics is another active research area that has recently shown how to identify an individual by analyzing their mouse movements in controlled settings [51, 64].

The construct of attention has nowadays become the common currency on the Web. Objective measurements of attentional processes are increasingly sought after by the media industry, to explain or predict user behavior. With every click or online interaction, digital footprints are created and logged, providing a detailed record of a person’s online activity that can be used for market segmentation, targeted advertising, but also for more privacy-invasive applications like user profiling.

Early mouse cursor tracking systems began by logging click events only (coordinates and timestamp) and using these events to assess what information users were interested in. However, it was soon realized that click data provide an incomplete picture of user interaction. Click data informed researchers of a users’ primary focus of attention, or their end choice. However, a mouse click is often preceded by several interactions such as scrolling, hovers, movements, etc. and thus can lead to a better overall understanding of the user’s thought process. This way, mouse cursor tracking systems began to incorporate such fine-grained within-page interactions to create richer user models.

In what follows, we review research efforts that have focused on mouse cursor analysis to infer user interest, visual attention, emotions, and demographic variables like gender or age, on a desktop setting. We thus deliberately leave out works on user profiling in mobile browsing, which fall outside the scope of this paper.

2.2.1 Inferring User Interest. For a long time, commercial search engines have been interested in how users interact with Search Engine Result Pages (SERPs), to anticipate better placement and allocation of ads in sponsored search or to optimize the content layout. Early work considered simple, coarse-grained features derived from mouse cursor data to be surrogate measurements of user interest [25, 87]. Follow-up research transitioned to more fine-grained mouse cursor features [32, 33] that were shown to be more effective. These approaches have been directed at predicting open-ended tasks like search success [36] or search satisfaction [63]. In a similar vein, Huang et al. [40, 41] modeled mouse cursor interactions and extended click models to compute more accurate relevance judgements of search results. Mouse cursor position is mostly aligned to eye gaze, especially on SERPs [34, 90], and that can be used as a good proxy for predicting good and bad abandonment [28].

2.2.2 Inferring Visual Attention. Mouse cursor tracking has been also used to survey the visual focus of users in sponsored search, thus revealing valuable – and at the same time sensitive – information regarding the distribution of user attention over the various SERP components. Despite the technical challenges that arise from this analysis, previous work has shown the utility of mouse movement patterns to measure within-content engagement [4] and predict reading experiences [5, 37]. Lagun et al. [52] introduced the concept of motifs, or frequent cursor subsequences, in the estimation of search result relevance. Similarly, Liu et al. [63] applied the motifs concept to SERPs and predicted search result utility, searcher effort, and satisfaction at a search task level. Boi et al. [16] proposed a method for predicting whether the user is looking at the content pointed by the cursor, exploiting the mouse cursor data and a segmentation of the contents in a web page. Lastly, Arapakis et al. [6, 8] investigated user engagement with direct displays on SERPs and provided further evidence that supports the utility of mouse cursor data for measuring user attention at a display-level granularity.

2.2.3 Inferring Emotional State. Although the connection between mouse cursor movements and the underlying psychological states has been a topic of research since the early 90s [3, 19], some studies have investigated the utility of mouse cursor data for predicting the user’s emotional state. For example, Zimmermann et al. [106] investigated the effect of induced affective states on the motor-behavior of online shoppers and found that the total duration of mouse cursor movements and the number of velocity changes were associated to the experienced arousal. Kaklauskas et al. [48] created a system that extracts physiological and motor-control parameters from mouse cursor interactions and then triangulated those with psychological data taken from self-reports, to analyse correlations with users’ emotional state and labour productivity. In a similar line, Azcarraga et al. [11] combined electroencephalography signals and mouse cursor interactions to predict self-reported emotions like frustration, interest, confidence and excitement. Yamauchi et

al. [103] studied the relationship between mouse cursor trajectories and generalized anxiety in human subjects. Lastly, Kapoor et al. [49] predicted whether a user experiences frustration, using an array of affective-aware sensors.

2.2.4 Inferring Demographics. Yamauchi et al. [104] examined the extent to which mouse cursor movements can help identify the gender and the experienced feelings of users who were watching short film clips. Although this work provides early evidence on the utility of mouse cursor data for advanced online user profiling, it suffers from certain limitations that we address in this work. First, the experimental setting has limited generalizability, since the adopted perception task is not very well connected to typical activities that users perform online, such as web search. Second, the data used in their predictive modeling task include multiple samples per participant randomly assigned to the training and test data partitions, hence there may be information leakage that artificially inflated model performance. In our analysis, we limit the training samples to exactly one mouse cursor trajectory per participant and test our models on unseen individuals.

Kratky et al. [50] recorded mouse cursor movements in an e-commerce website and engineered a set of meta-features to predict the user gender and age group. Their classifier was trained on several days of data per participant. Although the training and test collections had disjoint sets of participants, it was stated that the reported results were overly optimistic since researchers could not verify their ground-truth data [50]. In contrast, as discussed later, our dataset was collected from high-quality crowdworkers so we are confident that the ground-truth information is correct.

In a similar vein, Pentel et al. [78] used data from six different external sources, including e.g., keystroke data and feedback questionnaires, and handcrafted features proposed in earlier works [25, 28, 87] to train predictive models that could identify the users’ age and gender. However, because their approach relies mainly on ad-hoc data, it is less scalable and more difficult to implement than the approach we propose in this paper, which takes as input *raw* mouse cursor data. Moreover, Pentel et al. reported optimistic performance scores, which may be due to information leakage across data partitions, and omit important classification metrics such as precision, recall, and AUC. To account for their modeling approach, as well as that proposed by Kratky et al. [50], we implement the same classifier and test it in our setting (see Section 4).

2.3 Summary

Websites can infer fine-grained information about the users by tracking their mouse cursor activity. Tracking where exactly on the page a user’s mouse cursor hovers or clicks provides a surrogate signal for gaze fixation, and therefore reveals the focus of attention, which can be used to learn the users’ *latent* interests. However, the research literature on mouse cursor tracking has pointed out far more advanced and creative use cases for this technology. The above studies demonstrate that certain cognitive and motor control mechanisms are embodied and reflected, to some extent, in our mouse cursor movements and online interactions. In other words, mouse cursor movements can disclose sensitive information that may be employed for advanced user profiling, such as the identification of demographics, personality traits, and browsing intent.

For brevity’s sake, in the remainder of this paper we will focus on predicting demographics (age and gender) from mouse cursor movements, but we argue that other, potentially sensitive information may remain vulnerable and could be exposed if appropriate mechanisms are put into place.

3 STUDY

We ran an online user study a few years ago that reproduced the setting of a *sponsored search* task [57]. In order to make this paper self-contained, we will describe here the data collection procedure in enough detail to allow reproducibility of our work, nevertheless the reader may consult our reference paper [57] for more details.

Sponsored search provides the necessary revenue streams to commercial web search engines⁴ and it is critical to the success of many websites [46]. Commercial web search engines resort to various tracking techniques to monitor their users’ search activity, including mouse cursor tracking [27, 39, 41], and use that information to offer item recommendations [90], targeted advertising [13], or simply sell it to third parties [69].

With this work, we critique the use of mouse cursor tracking technology, highlighting possible implications for the future of the online advertising industry. More specifically, our user study allowed us to capture in a non-invasive manner the mouse cursor interactions of users who performed simple web search tasks. The collected mouse cursor data was then used to benchmark state-of-the-art machine learning models’ capacity to infer users’ demographic attributes.

3.1 Design

Our experiment, which was approved by a team of legal experts, consisted of a brief transactional search task [17] that was completed once per participant. Participants were presented with a simulated information need that explained that they were interested in purchasing a present for them or a friend, and were asked to use Google Search to find something appealing. Each participant was provided with a predefined search query and the corresponding SERP (see Figure 1) and were asked to click on any element of the SERP that answered it best. This way, we ensured that participants interacted with the same pool of web search queries and avoided any unaccounted systematic bias due to query quality variation. Overall, the search task consisted of three parts: (1) pre-task guidelines, (2) the web search task and (3) a post-task questionnaire.

The search queries (Section 3.2.1), which were all picked from a pool of popular queries in Google Search, were randomly distributed among our participants. The corresponding SERPs appeared all in English and were scraped for later instrumentation, simulating thus a website owner who wishes to track their users’ every move.

Participants accessed the instrumented SERPs through a dedicated server that did not alter the look and feel of the original SERPs. This allowed us to capture fine-grained user interactions while ensuring that the content of the SERPs remained consistent with the original version. Each participant was allowed to perform the search task *only once* to avoid introducing possible carry over effects and, thus, altering their browsing behavior in subsequent search tasks.

⁴<https://searchengineland.com/google-search-ad-revenues-271188>

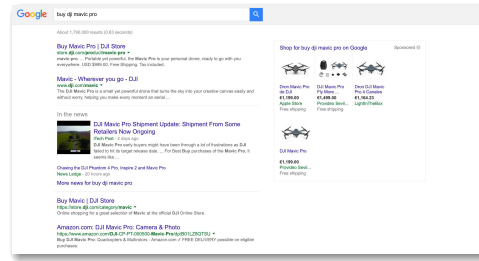


Figure 1: Example of a Google SERP used in the online study.

3.2 Apparatus

3.2.1 Search Query Sample. Starting from Google Trends,⁵ we selected a subset of the Top Categories and Shopping Categories that were suitable representatives of transactional tasks [17] i.e. categories that broadly express the intent of performing some web-mediated activity or transaction, like shopping or finding a service. Then, we extracted the top search queries issued in the US during the last 12 months and further narrowed down our search query collection to the 150 most popular search queries. Using this final selection of search queries, we produced the static version of the corresponding Google SERPs and injected JavaScript code (see next section) that allowed us to capture all client-side user interactions.

3.2.2 Mouse Cursor Tracking. As previously stated, all SERPs were downloaded and instrumented with custom JavaScript code. This way, we could automatically insert mouse tracking code and log cursor movements, hovers, and associated metadata. For this, we used EvTRACK,⁶ an open source JavaScript event tracking library derived from the smt2ε mouse tracking system [59].

We captured mousemove events via event polling, every 150 ms and all the other browser events (e.g., load, click, scroll) via event listeners. Whenever an event was recorded, we logged the following information: mouse cursor position (x and y coordinates), timestamp, event name, XPath of the DOM element that relates to the event, and the DOM element attributes (if any). EvTRACK has no dependencies and works in every major browser so, upon download, it is ready to use; i.e. no tooling or build pipeline is needed. This ease of use reveals how straightforward is to add mouse tracking capabilities to websites.

3.2.3 Questionnaire. In addition to the mouse cursor data, we gathered ground-truth information about the users through an online questionnaire that was administered at post-task. The questions included in the questionnaire were forced-choice type and allowed multi-point response options. The questionnaire comprised the following questions:

- (1) *What is your gender?* [Male, Female, Prefer not to say]
- (2) *What is your age group?* [18–23, 24–29, ..., 60–65, +66, Prefer not to say]
- (3) *What is your native language?* [Pull-down list, Prefer not to say]

⁵<https://trends.google.com/trends/>

⁶<https://github.com/luileito/evtrack>

3.3 Participants

We recruited participants from the FIGURE EIGHT crowdsourcing platform.⁷ They were of mixed nationality and had diverse educational backgrounds. All participants were proficient in English and were experienced (Level 3) contributors, i.e. they had a track record of successfully completed tasks and of a different variety, thus being considered very reliable contributors.

3.4 Procedure

Participants were instructed to read carefully the terms and conditions of the study which, among other things, informed them that they should perform the task from a desktop or laptop computer using a computer mouse (and refrain from using a touchpad, tablet, or mobile device) and that their browsing activity would be logged. Moreover, participants consented to share their browsing data and their questionnaire responses for later analysis.

Participants were asked to act naturally and choose anything that would best answer a given search query, since all “clickable” elements (e.g., result links, images, etc.) on the SERP were considered valid answers. The instructions were followed by a brief search task description like “*You want to buy <noun> (for you or someone else as a gift) and you have submitted the search query <noun> to Google Search. Please browse the search results page and click on the element that you would normally select under this scenario.*”

The SERPs were randomly assigned to the participants and each participant could take the study only once (see ‘Design’ section). Participants were allowed as much time as they needed to examine the SERP and proceed with the search task, which concluded upon selecting any of the “clickable” elements on the SERP. At the end of the task, participants were asked to complete the post-task questionnaire. The payment for participation was \$0.20 and the study took 0.83 minutes on average to complete (Mdn=0.37, SD=2.3) which roughly amounts to a 14\$/h wage. Participants could also opt-out at any moment, in which case they were not compensated.

3.5 Dataset

After excluding the users who did not provide their demographic information (see ‘Questionnaire’ section) and had few mouse movement data (less than ten mouse coordinates, which corresponds roughly to two seconds of user interaction data), we concluded on a set of 1,467 search sessions. The average mouse cursor trajectory length was 25.2 coordinates (SD=18.7, min=11, max=221). Next, our dataset was divided into a 90:10 training-test split; i.e., 90% of the data is used for model training and the remaining 10% of the data is used for testing. Our raw dataset is publicly available.⁸

3.6 Machine Learning Models

The focus of these experiments is demonstrating how feasible it is to implement a user profiling mechanism by relying on current machine learning techniques and easily acquired mouse cursor data. Therefore, for the sake of simplicity, we assume gender and age classification to be a two-class problem, i.e. a user is classified as ‘male’ or ‘female’ and as ‘young’ or ‘adult’. We note that age could be framed as a regression problem, however marketing companies

care more about market segmentation (i.e. fitting customers into target groups) rather than predicting a particular age [74, 86].

3.6.1 Baseline Models. We replicate the random forest (RF) classifier proposed in recent work [78, 104], which is an effective ensemble method that allows for a reliable performance assessment. Furthermore, we engineer a series of features (e.g., speed, acceleration, angle, traversed distance, hovers, clicks) and aggregate functions (e.g., min, max, mean and standard deviation) derived from the mouse cursor data, as reported in previous work [78, 104] (170 features). Then, we exclude the highly correlated ($r \geq .80$, $p < .05$) and linearly dependent features from our feature set and normalize the values for all features in the $[0, 1]$ range, so that feature values that fall in greater numeric ranges would not dominate those in smaller numeric ranges. In total, the RF model uses 52 mouse cursor features for classification. As a last step, we determine via grid search on a held-out set (comprising 10% of the training data) the optimal hyper-parameter values (number of trees, number of features, ϵ -threshold, minimum size of terminal nodes, maximum number of terminal nodes) and evaluate the performance of the RF model against the test set.

We also implement a ZeroR classifier, also known as 0-R (zero rule), which simply predicts the majority class. The ZeroR will always output the same target value and does not use any input features, hence its name. Despite its simplicity and lack of discriminative power, this classifier is very useful for determining the baseline performance, as a benchmark for other classification methods like the ones we used in these experiments.

3.6.2 Recurrent Neural Network Models. Creating a competent feature-based classifier like the RF previously described, as noted, demands significant effort and time because of the hyperparameter fine-tuning and, above all, the feature engineering process. Indeed, feature engineering requires domain expertise to derive features with sufficient discriminative power. With neural networks, however, feature engineering is automatically performed by the network itself. Together with the availability of state-of-the-art deep learning libraries such as Tensorflow, Keras, PyTorch, or MXNet, it has become increasingly easy to implement a competent classifier with few lines of code (see our implementation in Figure 6) and, hence, the purpose of this paper.

Since mouse movements are of sequential nature, we test a particular type of recurrent neural networks (RNNs) that is effective at modeling time series, where each data point in the sequence can be assumed to be dependent on the previous one. Concretely, the model uses Gated Recurrent Unit (GRU) memory, which is a simplification of the popular long short-term memory. We use the bidirectional variant (BiGRU) since a major issue with all RNNs is that they can only learn representations from *previous* time steps. However, sometimes we have to learn representations from *future* time steps to better understand the context and thus eliminate potential ambiguities.

Our BiGRU takes as input a *raw* sequence of mouse cursor positions and time offsets, which can be seen as a multivariate time series of three-dimensional data points. Because each mouse sequence has a different length, all sequences are padded to a fixed length of 100 timesteps, which corresponds roughly to the mean

⁷<https://www.figure-eight.com>

⁸<https://gitlab.com/iarapakis/the-attentive-cursor-dataset>




Figure 2: Age classification results. All metrics are weighted by class distribution.




Figure 3: Gender classification results. All metrics are weighted by class distribution.

sequence length observed in our dataset plus three standard deviations. The input layer of our RNN model has 100 neurons (one neuron per timestep). The hidden layer is the forward-backward recurrent block (BiGRU) with 64 output units, using hyperbolic tangent activation and sigmoid activation in the recurrent step. We add a dropout layer with drop rate $q = 0.25$ for regularization, followed by a fully-connected (FC) layer of 1 output unit using sigmoid activation. The model outputs a probability prediction p of the user’s gender or age, where $p > .5$ indicates that the user belongs to the majority class (in our data, ‘male’ and ‘young’ are the majority classes).

We use the popular Adam optimizer (stochastic gradient descent with momentum) with learning rate $\eta = 0.0005$ and decay rates $\beta_1 = 0.9$ and $\beta_2 = 0.999$. This model, including all the settings described above, is implemented in five lines of Python code; see Figure 6. We train this model with a batch size of 32 sequences and up to 400 epochs, with early stopping of 40 epochs to prevent overfitting.

4 USER PROFILING EXPERIMENTS

We report the weighted Precision, Recall, and F-measure (F1 score), according to the target class distributions in each case. In addition, we provide the Area Under the ROC curve (AUC), to highlight the discriminative power of each classifier. Finally, we remind the reader that the focus of this paper is not on attaining state-of-the-art performance but rather on demonstrating that it is feasible to implement a fairly competent user profiling mechanism by relying on current machine learning techniques and easily acquired mouse cursor data.

4.0.1 Age Classification. Prior work has linked age with motor control and pointing performance in tasks that involve the use of a computer mouse [15, 38, 47, 62, 88, 98]. Overall, ageing is marked by a decline in motor control abilities, therefore it is expected to affect the users’ pointing performance and, by extension, how they move the computer mouse. For example, Smith et al. [88] observed that older people incurred in longer mouse movement times, more sub-movements, and more pointing errors than the young. These

findings underline potential age effects on the way a mouse device is used in an online search task.

Figure 2 shows the performance results for the classification task that targets user age. Here, we divide our users into two age groups (“18–35” and “36–66”), in line with previous work [50, 78] that applied a comparable binary split on their user sample. While the RF model achieved an F-measure of 0.531 and an AUC of 0.528, the BiGRU outperformed its peers with an F-measure of 0.653 and an AUC of 0.712. Furthermore, we ran pairwise comparisons of proportions (Bonferroni-Holm corrected, to guard against over-testing the data) and observed statistically significant differences for all metrics when comparing the BiGRU against the other classifiers ($p < .01$).

We also note here that the performance of the RF model is much smaller than what researchers have reported in previous work [50, 78], whereas the simple implementation of our BiGRU model, only with raw mouse movements as input (spatial coordinates and time offsets) and five lines of code (Figure 6), validates the need to raise further awareness about the potential threats of mouse cursor tracking to online privacy.

4.0.2 Gender Classification. Prior research noted sensory-motor differences due to gender [23, 53, 105], such as significant variation in the cursor movement distance, pointing time, and cursor patterns. The cause of these variations has been attributed to gender-based differences in how users move a mouse cursor or to different cognitive mechanisms (perceptual and spatial processes) involved in motor control.

Figure 3 shows the performance results for the classification task that targets user gender. Again, the BiGRU model outperformed its peers. More specifically, the RF model achieved an F-measure of 0.523 and an AUC of 0.489, while the BiGRU achieved an F-measure of 0.641 and an AUC of 0.650. The pairwise comparisons of proportions (Bonferroni-Holm corrected) revealed statistically significant differences for all metrics except Recall when comparing the BiGRU against the other classifiers ($p < .01$). Although these results might not be as impressive as those pertaining age classification, they

clearly deviate from random classification and definitely call for attention to the potential implications for e-privacy.

In addition, we observe that, unlike the optimistic results reported by others [78, 104], the same RF model performed worse on our data, possibly due to the more challenging nature of the task. More importantly, we have shown that a shallow BiGRU model can outperform a predictive model that relies on a barrage of elaborate features, by using exclusively as input unprocessed mouse cursor movements, which are easy to acquire unobtrusively and at scale. Hence, sensitive information may be exploited by anyone who has access to a simple profiling technology, such as the one demonstrated in this paper.

5 PROFILING PREVENTION EXPERIMENTS

Now that we know that it is possible to easily infer user demographics from mouse cursor movements, we propose an adversarial method to modify the user’s movements in such a way that the resulting trajectory cannot disclose age and gender information. The method is illustrated in Figure 4: Whenever a `mousemove` event e_t happens at time t , we insert another `mousemove` event programmatically $e'_t \sim \mathcal{N}(0, \sigma)$ which is within a σ radius away from the original coordinate. This additive Gaussian noise is also applied to the time offsets, to ensure that the distorted trajectory has no duplicated times.




Figure 4: Adversarial noise example. We add an intermediate coordinate programmatically between two consecutive coordinates that is σ px away from the current position.

The amount of adversarial noise applied to each programmatic event ranges randomly from 0 to σ . We ensure that distorted points (both coordinates and time) are always positive values, in line with regular mouse movement data. In this experiment, we study the impact of σ in classification performance. Theoretically, a random classifier should achieve an AUC score of 0.5 for a two-class classification problem. Therefore, we expect to see a degradation in classification performance with regard to the previous experiments. Given that the BiGRU outperformed the RF model and relies only on raw mouse movements, we only challenge the BiGRU model in these experiments.

As observed in Figure 5, using a radius of $\sigma = 0.25$ px is enough to degrade the performance of the BiGRU model, which begins to behave as a random classifier ($AUC \approx 0.5$) of both age and gender. The differences between the original mouse data and the degraded versions are statistically significant at the $p < .001$ level. Hence, this




Figure 5: Degradation of classification results (weighted by class distribution) after introducing adversarial noise.

experiment justifies further our proposal to prevent user profiling techniques that exploit mouse cursor movements data. We argue that by using this adversarial noise, a mouse movement trajectory would become “illegible” to any machine learning model trying to classify the user’s age or gender. This experiment is inline with previous research that reported very small perturbations to cause a significant performance degradation [81]. Indeed, by scrambling both spatial and temporal information from a mouse cursor trajectory we are effectively signaling a seemingly arbitrary and jittery movement.

To validate further the validity of this adversarial noise, we re-trained our BiGRU model with distorted mouse data, using the same configuration from our previous experiments (Section 3.6.2). Now each mouse sequence is distorted according to a random uniform distribution $\sigma \sim \mathcal{U}(0, 1)$, which means that some coordinates are preserved ($\sigma = 0$) whereas others are more distorted ($\sigma = 1$). The results are shown in Table 1. As can be observed, the model achieves worse performance than the model trained on the original, non-distorted data, sometimes by a large margin. This was especially so for the age classifier. We conclude that the proposed adversarial noise technique is a robust countermeasure against the neural net used as profiling mechanism via mouse cursor tracking.

We have implemented this adversarial method in a Chrome extension (see ‘Resources’ section) that allows to configure the adversarial noise level (random uniform by default) and the number of `mousemove` events (one coordinate by default) to be added programmatically. The idea of providing these “sane defaults” is to avoid making the mouse movements too distorted so as they can go unnoticed by a machine learner trying to distinguish between human and fake movements.

6 DISCUSSION

The rapid growth of online advertising has spurred the demand for effective, but also at times privacy-invasive user profiling technologies that allow to deliver more relevant ad content to the user. Of course, user profiling is not a bad thing per se, since it allows to deliver more relevant ad content to the user. However, as targeted ads are believed to produce increased revenues, various intermediary companies (such as ad platforms and ad exchanges) are tracking users at scale, and often in an unregulated manner, which has resulted in a privacy nightmare [77]. Advertising is currently the

Demographics	Adj. Precision	Adj. Recall	Adj. F-measure	ROC AUC
Age	0.6301 ↓8.55%	0.5986 ↓11.05%	0.5328 ↓18.4%	0.6074 ↓14.7%
Gender	0.6429 ↑0.45%	0.6463 ↓0.04%	0.6133 ↓4.32%	0.6308 ↓2.95%

Table 1: Test performance of our BiGRU when trained using adversarial noise $\sigma \sim \mathcal{U}(0, 1)$. We show the performance degradation in parentheses, as a the percentage variation w.r.t the same classifier trained on non-distorted data.

main business model of “free” content and services on the Web, though if something is free it usually means that the user is the product. This dystopian reality evokes a disconcerting dichotomy between, on the one hand, having to accept a digital life with no privacy and, on the other hand, retaining our privacy by being off the digital life.

Privacy is essential for the citizens of both the physical and the digital world. But also privacy is constantly being juxtaposed with competing goods and interests, balanced against disparate needs and demands [80]. More importantly, the loss of privacy translates into a loss of freedom. In other words, freedom of expression is threatened by the surveillance of our digital traces; thought patterns and intentions can be extrapolated from website visits (rightly or wrongly), and the knowledge that we are being surveyed can make one less likely to research a particular topic.⁹ And even efforts to regulate the web tracking landscape, such as the DNT Header, require good faith cooperation from the parties at the other end of the web connection, which is not always guaranteed. In fact, most tracking domains and ad platforms are unlikely to ignore user tracking, because they make their business out of these data. Hence, the Web privacy problem is a fallout of rapid and uncontrolled growth in technology, mainly driven by a lack of transparency, control, and difficulty to understand e-privacy implications.

6.1 Implications and Outlook

Mouse cursor tracking is very difficult to avoid while browsing the Web today. Our mouse movements can be tracked silently at scale, in Incognito mode, and even without JavaScript enabled [42]. Being a low-cost and reasonable proxy of visual attention, mouse cursor tracking cannot be discounted from being a modern “Trojan horse”. Our analysis corroborates this account and demonstrates that, through a very simple machine learning implementation, we can infer people’s gender or age inadvertently. We do not argue, however, that mouse tracking should be removed from any website, as it may be a valuable data source for various application scenarios. Rather, we find it disconcerting that currently there is no way for end-users to opt-out easily.

The work presented here serves a dual purpose. First, it aims at raising awareness on the emerging privacy threats in the online world and exposes some of the unaccounted —yet sizeable risks— of tracking technologies. Even when dealing with a seemingly harmless browser activity logging practice, such as the collection of mouse cursor coordinates, a third party can mine this data to uncover personally identifiable information about the users. Second, with this work we intend to give control back to the users over their (mouse) data. We are certainly not the first ones in aiming

at this goal, though, to the best of our knowledge, our proposed adversarial noise technique is the first countermeasure against user profiling based on mouse cursor tracking. We note, however, that our method will not prevent “any and all” profiling techniques within mouse tracking, as the field of adversarial machine learning progresses rapidly and new counter-countermeasures are likely to appear in the future.

Researchers have proposed restricting access to only the features which are necessary for delivering a desired functionality [89], enforcing thus a principle of least privilege [26]. Google Chrome have recently announced a new technical proposal named “privacy budget” that could restore the balance between user privacy and ad targeting on the Web.¹⁰ With a privacy budget, websites could call APIs until those calls have revealed enough information, to narrow a user down to a group sufficiently large enough to maintain anonymity. After that, any further attempts to call such APIs would cause the browser to intervene and block further calls. Under this scenario, we could imagine a user-configurable privacy budget for mouse tracking data, though a set of sane defaults should be provided by browser vendors. For example, if a JavaScript function is listening to the `onmousemove` event more than N seconds in a row, the browser would block further calls to the listener function. We would also recommend browser vendors to list what sensitive APIs a website is using, just like they do currently to inform about SSL certificates, for example, or even ask for explicit consent to the user when the website requires access to such sensitive APIs, similar to app permission requests in mobile devices.

Finally, we hope that this work will motivate further research to counterbalance initiatives on developing privacy-invasive user profiling technologies by delivering techniques that can preserve user anonymity and protect personally identifiable information. This work also highlights the need for more transparency and privacy-aware tools. We believe that users should be tracked, if at all, by category instead of individually. While some advertisers do care about organizing users into general groups, others aim at creating detailed individual profiles, which should not happen without explicit user’s consent.

6.2 Limitations and Future Work

We have analyzed movements generated by a computer mouse and so the proposed method is not expected to work “as is” on touch-capable devices, such as tablets or smartphones. However, user engagement is still higher on desktop than on mobile [8], which means more profitability for advertisers. Nonetheless, this presents an opportunity to extend our work and account for touch-based

⁹<https://robindaherty.com/2016/01/06/nothing-to-hide.html>

¹⁰<https://blog.chromium.org/2019/08/potential-uses-for-privacy-sandbox.html>

interactions such as, for example, tracking zoom/pinch gestures and scroll activity instead of the mouse cursor position [35].

We have shown that it is possible to detect user demographics with reasonable accuracy. More importantly, we have shown that it is possible to do so unobtrusively and at scale, by relying only on sequences of raw mouse cursor data. However, since the focus of our work is not on attaining state-of-the-art performance, there is still room to benchmark further the capabilities of such user profiling technologies and uncover additional vulnerabilities in the data. For example, stacking more recurrent layers (deeper model), increasing the number of hidden neurons (wider model), or using data augmentation techniques. Even non-sequential models are also possible to analyze mouse cursor data [7].

Finally, we acknowledge a limitation of our adversarial noise technique. The W3C consortium introduced the concept of “trusted events”,¹¹ to help developers differentiate between events triggered by a genuine user interaction and those triggered programmatically, e.g., by a 3rd party script. Our Chrome extension adds mouse cursor distortions programmatically via JavaScript, therefore those events are considered untrusted, although currently none of the major mouse tracking companies filter out untrusted DOM events.¹² It is a matter of time, however, for companies to catch up and update their tracking technology. Therefore, in future work we will release a program that runs at the Operating System level¹³ and thus can trigger mouse events that are seen as trusted by the web browser.

7 CONCLUSION

It is possible to infer user demographics unobtrusively and at scale with reasonable accuracy, using an off-the-shelf recurrent neural network that takes as input raw mouse movements. Previous attempts have relied on expert knowledge in machine learning techniques and feature engineering methods. Therefore, we noticed an unprecedented low entry barrier for webmasters interested in profiling the user on their websites with no effort, highlighting thus a disconcerting privacy issue. We have proposed an adversarial noise method to mitigate such user profiling techniques that make use of mouse cursor tracking to predict demographic variables such as gender and age, so that users interested in preserving their privacy can do so with no effort too. It is our hope that this paper will raise awareness among the research community about how easy the task of profiling users on the Web has become, including mouse cursor tracking and beyond. With this paper we want to bring together browser vendors, advertising platforms, practitioners, and web users to reflect on the tradeoffs between privacy and technological innovation, and the impact that unethical practices may have on users in the real world.

ACKNOWLEDGMENTS

We thank Heather O’Brien and Irene Lopatovska for providing early feedback, as well as our anonymous reviewers.

¹¹<https://www.w3.org/TR/uitablevents/#trusted-events>

¹²See e.g., <https://easylist.to/easylist/easyprivacy.txt>

¹³<https://github.com/jordansissel/xdotool>

A RNN CODE

Figure 6 illustrates the deceptively ease of creating a fairly competent recurrent neural network for two-class classification that takes as input a trajectory of (x, y, t) coordinates (max. 100 timesteps) and outputs the majority class probability.

```
1 | model = Sequential()
2 | model.add(Bidirectional(GRU(64), input_shape = (100, 3)))
3 | model.add(Dropout(0.25))
4 | model.add(Dense(1, activation = "sigmoid"))
5 | model.compile(loss = "binary_crossentropy", optimizer =
   |     Adam(0.0005))
```

Figure 6: Our RNN implementation with the Keras library.

B RESOURCES

We release a Chrome extension that implements our adversarial noise approach to distort the mouse cursor coordinates on the fly. The extension can be enabled or disabled for whitelisted domains. This way, the user can allow certain websites to track their mouse movements as needed; e.g., as part of an auditing process of a banking website, an e-commerce that do not request personal data but want to get a demographics overview of their visitors, or simply a research study that pays the user for letting them to analyze their mouse cursor activity. The extension can be downloaded at <https://github.com/luileito/mousefaker>.

Our dataset comprising mouse cursor movements and associated demographic variables is available at <https://gitlab.com/iarapakis/the-attentive-cursor-dataset>. Each user log includes the following information: query, gender, age, browser viewport size (width and height), and mouse cursor trajectory as a sequence of (x, y, t) tuples.

REFERENCES

- [1] G. Acar, C. Eubank, S. Englehardt, M. Juarez, A. Narayanan, and C. Diaz. 2014. The Web Never Forgets: Persistent Tracking Mechanisms in the Wild. In *Proc. CCS*.
- [2] G. Acar, M. Juarez, N. Nikiforakis, C. Diaz, S. Gürses, F. Piessens, and B. Preneel. 2013. FPDetective: Dusting the Web for Fingerprinters. In *Proc. CCS*.
- [3] J. Accot and S. Zhai. 1997. Beyond Fitts’ Law: Models for Trajectory-based HCI Tasks. In *Proc. CHI*.
- [4] I. Arapakis, M. Lalmas, B. B. Cambazoglu, M.-C. Marcos, and J. M. Jose. 2014. User Engagement in Online News: Under the Scope of Sentiment, Interest, Affect, and Gaze. *J. Assoc. Inf. Sci. Technol.* 65, 10 (2014).
- [5] I. Arapakis, M. Lalmas, and G. Valkanas. 2014. Understanding Within-Content Engagement Through Pattern Analysis of Mouse Gestures. In *Proc. CIKM*.
- [6] I. Arapakis and L. A. Leiva. 2016. Predicting User Engagement with Direct Displays Using Mouse Cursor Information. In *Proc. SIGIR*.
- [7] I. Arapakis and L. A. Leiva. 2020. Learning Efficient Representations of Mouse Movements to Predict User Attention. In *Proc. SIGIR*.
- [8] I. Arapakis, A. Penta, H. Joho, and L. A. Leiva. 2020. A Price-Per-Attention Auction Scheme Using Mouse Cursor Information. *ACM Trans. Inf. Syst.* 38, 2 (2020).
- [9] E. Arroyo, T. Selker, and W. Wei. 2006. Usability tool for analysis of web designs using mouse tracks. In *Proc. CHI EA*.
- [10] R. Atterer, M. Wnuk, and A. Schmidt. 2006. Knowing the User’s Every Move: User Activity Tracking for Website Usability Evaluation and Implicit Interaction. In *Proc. WWW*.
- [11] J. Azcarraga and M. T. Suarez. 2012. Predicting Academic Emotions Based on Brainwaves, Mouse Behaviour and Personality Profile. In *Proc. PRICAI*.
- [12] P. Bangera and S. Gorinsky. 2017. Ads versus Regular Contents: Dissecting the Web Hosting Ecosystem. In *Proc. IFIP Networking*.
- [13] M. Bashir, S. Arshad, W. Robertson, and C. Wilson. 2016. Tracing Information Flows Between Ad Exchanges Using Retargeted Ads. In *Proc. SEC*.
- [14] H. Beals. 2010. The value of behavioral targeting. Available at <https://www.networkadvertising.org>.

- [15] M. Bohan and A. Chaparro. 1998. Age-Related Differences in Performance Using a Mouse and Trackball. *Hum. Factors* 42, 2 (1998).
- [16] P. Boi, G. Fenu, L. D. Spano, and V. Vargiu. 2016. Reconstructing User's Attention on the Web Through Mouse Movements and Perception-Based Content Identification. *ACM Trans. Appl. Percept.* 13, 3 (2016).
- [17] A. Broder. 2002. A Taxonomy of Web Search. *SIGIR Forum* 36, 2 (Sept. 2002), 3–10. <https://doi.org/10.1145/792550.792552>
- [18] T. I. A. Bureau. 2019. Standards, Guidelines and Best Practices. Available at <https://www.iab.com>.
- [19] S. K. Card, W. K. English, and B. J. Burr. 1987. Evaluation of Mouse, Rate-controlled Isometric Joystick, Step Keys, and Text Keys, for Text Selection on a CRT. In *Human-computer Interaction*, R. M. Baecker and W. A. S. Buxton (Eds.). Taylor & Francis.
- [20] J. P. Carrascal, C. Riederer, V. Erramilli, M. Cherubini, and R. de Oliveira. 2013. Your Browsing Behavior for a Big Mac: Economics of Personal Information Online. In *Proc. WWW*.
- [21] J. M. Carrascosa, J. Mikians, R. Cuevas, V. Erramilli, and N. Laoutaris. 2015. I Always Feel Like Somebody's Watching Me: Measuring Online Behavioural Advertising. In *Proc. CoNEXT*.
- [22] M. C. Chen, J. R. Anderson, and M. H. Sohn. 2001. What Can a Mouse Cursor Tell Us More? Correlation of Eye/Mouse Movements on Web Browsing. In *Proc. CHI EA*.
- [23] R. C. C. Chen and T.-K. Chen. 2008. The effect of gender-related difference on human-centred performance using a mass assessment method. *IJCAT* 32 (2008).
- [24] Y. Chen, Y. Liu, M. Zhang, and S. Ma. 2017. User Satisfaction Prediction with Mouse Movement Information in Heterogeneous Search Environment. *IEEE Trans. Knowl. Data. Eng.* 29, 11 (2017).
- [25] M. Claypool, P. Le, M. Wased, and D. Brown. 2001. Implicit Interest Indicators. In *Proc. IUI*.
- [26] P. J. Denning. 1976. Fault tolerant operating systems. *ACM Comput. Surv.* 8, 4 (1976).
- [27] F. Diaz, R. White, G. Buscher, and D. Liebling. 2013. Robust Models of Mouse Movement on Dynamic Web Search Results pages. In *Proc. CIKM*.
- [28] A. Diriyee, R. White, G. Buscher, and S. Dumais. 2012. Leaving So Soon? Understanding and Predicting Web Search Abandonment Rationales. In *Proc. CIKM*.
- [29] S. Englehardt. 2017. No boundaries: Exfiltration of personal data by session-replay scripts. Available at <https://freedom-to-tinker.com>.
- [30] H. A. Feild, J. Allan, and R. Jones. 2010. Predicting Searcher Frustration. In *Proc. SIGIR*.
- [31] K. Z. Gajos, K. Reinecke, M. Donovan, C. D. Stephen, A. Y. Hung, J. D. Schmahmann, and A. S. Gupta. 2020. Computer mouse use captures ataxia and parkinsonism, enabling accurate measurement and detection. *Mov. Disord.* 35, 2 (2020).
- [32] Q. Guo and E. Agichtein. 2008. Exploring Mouse Movements for Inferring Query Intent. In *Proc. SIGIR*.
- [33] Q. Guo and E. Agichtein. 2010. Ready to Buy or Just Browsing? Detecting Web Searcher Goals from Interaction Data. In *Proc. SIGIR*.
- [34] Q. Guo and E. Agichtein. 2012. Beyond Dwell Time: Estimating Document Relevance from Cursor Movements and Other Post-click Searcher Behavior. In *Proc. WWW*.
- [35] Q. Guo, H. Jin, D. Lagun, S. Yuan, and E. Agichtein. 2013. Mining Touch Interaction Data on Mobile Devices to Predict Web Search Result Relevance. In *Proc. SIGIR*.
- [36] Q. Guo, D. Lagun, and E. Agichtein. 2012. Predicting Web Search Success with Fine-grained Interaction Data. In *Proc. CIKM*.
- [37] D. Hauger, A. Paramythis, and S. Weibelzahl. 2011. Using browser interaction data to determine page reading behavior. In *Proc. UMAP*.
- [38] S. H. Hsu, C. C. Huang, Y. H. Tsuang, and J. S. Sun. 1999. Effects of age and gender on remote pointing performance and their design implications. *Int. J. Ind. Ergon.* 23, 5 (1999).
- [39] J. Huang, R. White, and G. Buscher. 2012. User See, User Point: Gaze and Cursor Alignment in Web Search. In *Proc. CHI*.
- [40] J. Huang, R. W. White, G. Buscher, and K. Wang. 2012. Improving Searcher Models Using Mouse Cursor Activity. In *Proc. SIGIR*.
- [41] J. Huang, R. W. White, and S. Dumais. 2011. No Clicks, No Problem: Using Cursor Movements to Understand and Improve Search. In *Proc. CHI*.
- [42] L.-S. Huang, Z. Weinberg, C. Evans, and C. Jackson. 2010. Protecting Browsers from Cross-origin CSS Attacks. In *Proc. CCS*.
- [43] C. Iordanou, G. Smaragdakis, and N. Laoutaris. 2019. Who's Tracking Sensitive Domains? *CoRR* abs/1908.02261 (2019).
- [44] C. Iordanou, G. Smaragdakis, I. Poese, and N. Laoutaris. 2018. Tracing Cross Border Web Tracking. In *Proc. IMC*.
- [45] C. Jackson, A. Bortz, D. Boneh, and J. C. Mitchell. 2006. Protecting Browser State from Web Privacy Attacks. In *Proc. WWW*.
- [46] B. J. Jansen and T. Mullen. 2008. Sponsored search: an overview of the concept, history, and technology. *Int. J. Electronic Business* 6 (2008).
- [47] T. Jastrzebski, N. Charness, P. Holley, and J. Feddon. 2003. Input devices for web browsing: age and hand effects. *Universal Access Inf.* 4 (2003).
- [48] A. Kaklauskas, M. Krutinis, and M. Seniut. 2009. Biometric Mouse Intelligent System for Student's Emotional and Examination Process Analysis. In *Proc. ICALT*.
- [49] A. Kapoor, W. Bursleson, and R. W. Picard. 2007. Automatic Prediction of Frustration. *Int. J. Hum.-Comput. Stud.* 65, 8 (2007).
- [50] P. Kratky and D. Chuda. 2016. Estimating Gender and Age of Web Page Visitors from the Way They Use Their Mouse. In *Proc. WWW Companion*.
- [51] P. Krátky and D. Chudá. 2018. Recognition of Web Users with the Aid of Biometric User Model. *J. Intell. Inf. Syst.* 51, 3 (2018).
- [52] D. Lagun, M. Ageev, Q. Guo, and E. Agichtein. 2014. Discovering Common Motifs in Cursor Movement Data for Improving Web Search. In *Proc. WSDM*.
- [53] A. A. Landauer. 1981. Sex Differences in Decision and Movement Time. *Percept. Mot. Skills* 52, 1 (1981).
- [54] N. Laoutaris. 2016. Cows, privacy, and tragedy of the commons on the Web. Available at <http://laoutaris.info>.
- [55] M. Lecuyer, G. Ducoffé, F. Lan, A. Papancea, T. Petsios, R. Spahn, A. Chain-treau, and R. Geambasu. 2014. XRay: Enhancing the Web's Transparency with Differential Correlation. In *Proc. SEC*.
- [56] L. A. Leiva. 2011. Restyling Website Design via Touch-based Interactions. In *Proc. MobileHCI*.
- [57] L. A. Leiva and I. Arapakis. 2020. The Attentive Cursor Dataset. *Front. Hum. Neurosci.* 14 (2020).
- [58] L. A. Leiva and J. Huang. 2015. Building a better mousetrap: Compressing mouse cursor activity for web analytics. *Inf. Process. Manag.* 51, 2 (2015).
- [59] L. A. Leiva and R. Vivó. 2013. Web Browsing Behavior Analysis and Interactive Hypervideo. *ACM Trans. Web* 7, 4 (2013).
- [60] A. Lerner, A. K. Simpson, T. Kohno, and F. Roesner. 2016. Internet Jones and the Raiders of the Lost Trackers: An Archaeological Study of Web Tracking from 1996 to 2016. In *Proc. SEC*.
- [61] C. Leung, J. Ren, D. Choffnes, and C. Wilson. 2016. Should You Use the App for That? Comparing the Privacy Implications of App- and Web-based Online Services. In *Proc. IMC*.
- [62] T. Lindberg, R. Näsänen, and K. Müller. 2006. How age affects the speed of perception of computer icons. *Displays* 27, 4 (2006).
- [63] Y. Liu, Y. Chen, J. Tang, J. Sun, M. Zhang, S. Ma, and X. Zhu. 2015. Different Users, Different Opinions: Predicting Search Satisfaction with Mouse Movement Information. In *Proc. SIGIR*.
- [64] H. Lu, J. Rose, Y. Liu, A. Awad, and L. Hou. 2017. Combining Mouse and Eye Movement Biometrics for User Authentication. In *Information Security Practices*, I. Traoré, A. Awad, and I. Woungang (Eds.). Springer.
- [65] D. Malandrino, V. Scarano, and R. Spinelli. 2013. How increased awareness can impact attitudes and behaviors toward online privacy protection. In *Proc. PASSAT*.
- [66] D. Martín-Albo, L. A. Leiva, J. Huang, and R. Plamondon. 2016. Strokes of insight: User intent detection and kinematic compression of mouse cursor trails. *Inf. Process. Manag.* 52, 6 (2016).
- [67] A. M. McDonald and L. F. Cranor. 2010. Americans' Attitudes About Internet Behavioral Advertising Practices. In *Proc. WPES*.
- [68] A. M. McDonald and L. F. Cranor. 2010. Beliefs and behaviors: Internet users' understanding of behavioral advertising. In *Proc. TPRC*.
- [69] T. Morey, T. Forbath, and A. Schoop. 2015. Customer Data: Designing for Transparency and Trust. Harvard Business Review.
- [70] K. Mowery, D. Bogenreif, S. Yilek, and H. Shacham. 2011. Fingerprinting Information in JavaScript Implementations. In *Proc. W2SP*.
- [71] F. Mueller and A. Lockerd. 2001. Cheese: Tracking Mouse Movement Activity on Websites, a Tool for User Modeling. In *Proc. CHI EA*.
- [72] V. Navalpakkam, L. Jentzsch, R. Sayres, S. Ravi, A. Ahmed, and A. Smola. 2013. Measurement and Modeling of Eye-mouse Behavior in the Presence of Nonlinear Page Layouts. In *Proc. WWW*.
- [73] S. of California. 2018. California Consumer Privacy Act – Assembly Bill No. 375. Available at <https://leginfo.ca.gov>.
- [74] H. Oh, S. C. Parks, and F. J. Demicco. 2002. Age- and Gender-Based Market Segmentation: A Structural Understanding. *Int. J. Hosp. Tour. Adm.* 3, 1 (2002).
- [75] L. Olejnik, C. Castelluccia, and A. Janc. 2012. Why Johnny Can't Browse in Peace: On the Uniqueness of Web Browsing History Patterns. In *Proc. HotPETS Workshop*.
- [76] P. Papadopoulos, P. Rodriguez, N. Kourtellis, and N. Laoutaris. 2017. If you are not paying for it, you are the product: how much do advertisers pay to reach you?. In *Proc. IMC*.
- [77] J. Parra-Armau, J. P. Achara, and C. Castelluccia. 2017. MyAdChoices: Bringing Transparency and Control to Online Advertising. *ACM Trans. Web* 11 (2017).
- [78] A. Pentel. 2017. Predicting Age and Gender by Keystroke Dynamics and Mouse Patterns. In *Adj. Proc. UMAP*.
- [79] A. C. Plane, E. M. Redmiles, M. M. Mazurek, and M. C. Tschantz. 2017. Exploring User Perceptions of Discrimination in Online Targeted Advertising. In *Proc. USENIX Security*.
- [80] D. E. Pozen. 2016. Privacy-Privacy Tradeoffs. *The University of Chicago Law Review* 83, 1 (2016).

- [81] A. Prakash, N. Moran, S. Garber, A. DiLillo, and J. Storer. 2018. Deflecting Adversarial Attacks with Pixel Deflection. In *Proc. CVPR*.
- [82] E. Pujol, O. Hohlfeld, and A. Feldmann. 2015. Annoyed Users: Ads and Ad-Block Usage in the Wild. In *Proc. IMC*.
- [83] A. Razaqhanah, R. Nithyanand, N. Vallina-Rodriguez, S. Sundaresan, M. Allman, C. Kreibich, and P. Gill. 2018. Apps, Trackers, Privacy, and Regulators: A Global Study of the Mobile Tracking Ecosystem. In *Proc. NDSS*.
- [84] J. Ren, A. Rao, M. Lindorfer, A. Legout, and D. Choffnes. 2016. ReCon: Revealing and Controlling PII Leaks in Mobile Network Traffic. In *Proc. MobiSys*.
- [85] F. Schaub, A. Marella, P. Kalvani, B. Ur, C. Pan, E. Forney, and L. F. Cranor. 2016. Watching Them Watching Me: Browser Extensions' Impact on User Privacy Awareness and Concern. In *Proc. USEC Workshop at NDSS*.
- [86] C. D. Schewe and S. M. Noble. 2000. Market Segmentation by Cohorts: The Value and Validity of Cohorts in America and Abroad. *J. Mark. Manag.* 16, 1–3 (2000).
- [87] B. Shapira, M. Taieb-Maimon, and A. Moskowit. 2006. Study of the Usefulness of Known and New Implicit Indicators and Their Optimal Combination for Accurate Inference of Users Interests. In *Proc. SAC*.
- [88] M. W. Smith, J. Sharit, and S. J. Czaja. 1999. Aging, Motor Control, and the Performance of Computer Mouse Tasks. *Hum. Factors* 41, 3 (1999).
- [89] P. Snyder, C. Taylor, and C. Kanich. 2017. Most Websites Don't Need to Vibrate: A Cost-Benefit Approach to Improving Browser Security. In *Proc. CCS*.
- [90] M. Speicher, A. Both, and M. Gaedke. 2013. TellMyRelevance! Predicting the Relevance of Web Search Results from Cursor Interactions. In *Proc. CIKM*.
- [91] O. Starov, P. Gill, and N. Nikiforakis. 2016. Are You Sure You Want to Contact Us? Quantifying the Leakage of PII via Website Contact Forms. In *Proc. PoPETS*.
- [92] E. Steven and A. Narayanan. 2016. Online Tracking: A 1-million-site Measurement and Analysis. In *Proc. CCS*.
- [93] L. Sweeney. 2013. Discrimination in online ad delivery. *Commun. ACM* 56, 5 (2013).
- [94] J. Y. Tsai, S. Egelman, L. Cranor, and A. Acquisti. 2011. The effect of online privacy information on purchasing behavior: An experimental study. *Inf. Sys. Res.* 22, 2 (2011).
- [95] J. Turow, J. King, C. J. Hoofnagle, A. Bleakley, and M. Hennessy. 2009. Americans Reject Tailored Advertising and Three Activities that Enable It. *SSRN Electronic Journal*.
- [96] E. Union. 2016. The EU General Data Protection Regulation. Available at <http://eur-lex.europa.eu>.
- [97] B. Ur, P. G. Leon, L. F. Cranor, R. Shay, and Y. Wang. 2012. Smart, useful, scary, creepy: Perceptions of online behavioral advertising. In *Proc. SOUPS*.
- [98] N. Walker, D. A. Philbin, and A. D. Fisk. 1997. Age-related differences in movement control: Adjusting submovement structure to optimize performance. *J. Gerontol. A Biol. Sci. Med. Sci.* 52, 1 (1997).
- [99] R. J. Walls, E. D. Kilmer, N. Lageman, and P. D. McDaniel. 2015. Measuring the Impact and Perception of Acceptable Advertisements. In *Proc. IMC*.
- [100] R. White, P. Doraiswamy, and E. Horvitz. 2018. Detecting neurodegenerative disorders from web search signals. *npj Digital Med.* 1, 8 (2018).
- [101] C. E. Wills and C. Tatar. 2012. Understanding What They Do with What They Know. In *Proc. WPES*.
- [102] C. E. Wills and M. Zeljkovic. 2010. A personalized approach to web privacy: Awareness, attitudes and actions. *Inform. Manag. Comput. Secur.* 19, 1 (2010).
- [103] T. Yamauchi. 2013. Mouse Trajectories and State Anxiety: Feature Selection with Random Forest. In *Proc. ACII*.
- [104] T. Yamauchi and C. Bowman. 2014. Mining Cursor Motions to Find the Gender, Experience, and Feelings of Computer Users. In *Proc. ICDMW*.
- [105] T. Yamauchi, J. H. Seo, N. Jett, G. Parks, and C. Bowman. 2015. Gender Differences in Mouse and Cursor Movements. *Int. J. Hum.-Comput. Interact.* 31, 12 (2015).
- [106] P. Zimmermann, S. Guttormsen, B. Danuser, and P. Gomez. 2003. Affective Computing – A Rationale for Measuring Mood With Mouse and Keyboard. *Int. J. Occup. Saf. Ergon.* 9 (2003).