

Research Statement

Hieu Le

levanhieu.com

Protecting user privacy and improving awareness of data collection practices by online platforms are major challenges. Personal data are routinely collected to personalize experiences and provide services. However, it is also used to manipulate buying behaviors (*e.g.*, ad targeting), spread misinformation (*e.g.*, political ads), and it can even be sold to undisclosed parties. My long-term career goal is to give people control over their personal data and experiences in online platforms. My research will advance this goal by developing frameworks and methodologies to enhance transparency, auditability, and control for end-users.

Overview

Transparency and Auditability. Personal data collection practices are often opaque to end-users and researchers, especially for emerging platforms. My research elucidates these practices by conducting large-scale network measurements of apps on smart TVs and Oculus VR, employing specialized techniques to decrypt network traffic, extracting data types that are exposed (*e.g.*, personal identifiable information (PII), metadata that can be utilized to fingerprint the user, and platform-specific data), and identifying the organizations that collect the data and for which purposes, such as advertising and tracking (A&T) [9, 10]. We find that first, third, and platform parties, all collect a multitude of data types, which makes identifying users possible. Furthermore, to address auditability, my research develops tools to compare these practices with their expected behavior, which are commonly disclosed to end-users through privacy policies [9]. We find that data is often collected for nonessential purposes (*e.g.*, marketing vs. functionality) and privacy policies are often vague about the data being collected. Our research and findings have been presented to the Federal Trade Commission (FTC) at PrivacyCon 2021 and 2022 [7, 6], Consumer Reports, and privacy-focused industry players such as Duck Duck Go. Some of our research methodologies have been integrated into workshops that educate underrepresented college students about IoT privacy [8].

Personal Data Control. Beyond informing end-users about how their data is collected, users should also have the ability to control whether their data is collected or not. Millions of users worldwide depend on privacy-enhancing technologies (PETs), such as adblockers, to block A&T. However, the efficacy of PETs often rely on trial-and-error human processes that are error-prone and do not scale well across millions of websites and apps — and especially, over time. My research fortifies the effectiveness of PETs for the web, by identifying their pain-points and developing frameworks and methodologies that automate them [4, 5]. In particular, we train a machine learning (ML) classifier to notify human experts when to update PETs for websites where it is no longer effective and develop a reinforcement learning (RL) framework and tool to automate those updates. Our findings show that human experts can update PETs almost hourly — our RL-based tool can reduce this human effort by automating this process with comparable performance to the human experts for the top-thousands of websites. Our research has been presented for the last four years at the Ad-Filtering Dev Summit, where companies that develop PETs (*e.g.*, browsers, adblockers) gather to disseminate emerging ideas for PETs and even how they could impact regulations [2]. Some of our research has been customized and adopted by these companies [3, 5].

Current Research

Robust and Scalable PETs for the Web. PETs, such as adblockers for the web, are powered by filter rules; they are string-based patterns that can match and block A&T requests. However, rules are manually created and maintained by human experts. Conversely, publishers and advertisers work with third-party services that provide specialized techniques to circumvent adblockers to recover revenue. In return, human experts must update filter rules as well. To understand this accelerated arms race, we conduct a longitudinal analysis to measure how often experts update filter rules by inspecting their Github history: they update rules almost hourly to combat circumvention. To reduce this human effort, we build an AI-based tool, CV-Inspector [5], that can detect when websites have circumvented adblockers. First, we identify the specialized techniques used to circumvent adblockers such as randomizing URL components (*e.g.*, subdomains, paths), obfuscating JavaScript code, and obfuscated HTML ad structures. We leverage these findings to extract features (from HTTP requests and HTML DOM modalities) based on differential analysis (*i.e.*, how a website naturally behaves *vs.* how it behaves when a user has an adblocker) to train an ML classifier. Our evaluation shows that our model can automatically detect when websites circumvent adblockers, effectively reducing the human effort by 98%. This work was presented at the Ad-Filtering Dev Summit and customized for industry use in 2021.

Although CV-Inspector can automatically notify human experts when to update filter rules, it does not help create rules. In response, researchers have built AI-based tools that automatically detect and block A&T or designed approaches that assist in the creation of rules. However, these approaches have two main limitations. First, they rely on existing filter rules to help label their ground truth, causing a circular dependency — experts now must maintain both filter rules and AI-based tools. Second, they do not automatically evaluate whether the decision to block a request will cause breakage to a website (*e.g.*, missing legitimate images and text). To address both limitations, we present AutoFR [4], a framework based on reinforcement learning (RL) and a practical tool, that can automatically generate URL-based filter rules from scratch for a given website, that block ads while avoiding visual breakage [4]. We formulate the problem using multi-arm bandits. The human is now a user that provides AutoFR with the website to generate rules for, and a threshold that denotes how much they care about avoiding breakage. Whereas, the RL agent is tasked with learning which rules are effective at blocking ads without causing visual breakage beyond the given threshold, and within a time limit. To implement a scalable tool, we represent how a site is loaded using a provenance-based graph of how resources such as JS code, images, and text are loaded. Now, instead of visiting a live site (which can be slow), we read the graph into memory and apply the rule to infer its effectiveness. Using this approach, it takes on average 1.6 minutes to generate rules per-site. We apply AutoFR on the Top-5K sites and find that it creates rules with comparable performance to the state-of-the-art filter list, EasyList. This work was presented at the Ad-Filtering Dev Summit in 2022.

Characterizing the Advertising and Tracking Ecosystems of Emerging Platforms. Although there is extensive work for the web and mobile, privacy on emerging platforms, such as smart TVs and VR headsets, are not well-understood. For smart TVs, their increasing prevalence in households worldwide provide new opportunities for A&T. For VR headsets, such as the Oculus VR by Facebook, they include new sensors that detect user facial features, movements, and environments, that can lead to exposures of sensitive information such as the user’s age, gender, and economic status. To that end, our research involves developing methodologies to measure data collection practices within these emerging platforms to improve transparency for end-users [9, 10].

We conduct large-scale network measurement studies of smart TVs and the Oculus VR by employing dynamic code instrumentation, binary analysis, and an on-device VPN, to collect the network traffic. We implement a semi-automatic approach to extracting data types within the network traffic using hardcoded

data within device settings and regular expressions. We find that both platforms collect a wide variety of PII (*e.g.*, email, serial number), metadata that can be used for fingerprinting (*e.g.*, device OS, Unity version), and platform-specific data (*e.g.*, VR movement, VR play area). To understand whether the data was collected for A&T, we apply state-of-the-art filter lists (that contain domains of A&T) to the network traffic. For smart TVs, the A&T ecosystem is diverse with organizations like Alphabet, Facebook, and comScore. Conversely, we observe that the A&T ecosystem for Oculus VR is still developing, with few companies, such as Unity and Facebook, that track users for social and analytics purposes. At the time, there were no on-device ads for the Oculus VR. Our research and findings have been presented to the FTC in 2021 and 2022, Consumer Reports, and Duck Duck Go.

Future Directions

My future research will continue along the themes of transparency, auditability, and control for end-users. I plan to expand my research scope to study privacy in emerging platforms, to foster human-AI interactions and trust for PETs, and to develop scalable methodologies to detect and control dark patterns.

1. Privacy in Emerging Platforms

Privacy Implications for Extended Reality Platforms. New platforms introduce immersive experiences for end-users such as those coming from extended reality (XR) devices (*e.g.*, augmented, virtual, and mixed reality). However, they include new sensors that can learn and expose sensitive information about users, such as the users' age, gender, and economic status. I will leverage my expertise from the web, smart TVs, and Oculus VR, to conduct research that studies which data types can be learned and inferred by XR platforms, who are collecting these data types, and for which purposes. Furthermore, immersive experiences can influence user behavior and emotions more easily. As a result, I plan to investigate how XR advertising can harm end-users (*e.g.*, manipulate user buying behavior, manipulate emotions to spread misinformation) and work with committees, such as the Acceptable Ads Committee [1], to mitigate these potential harms by designing XR ad standards.

Robust and Scalable PETs for Other Platforms. Mobile, smart TVs, and XR platforms are increasingly becoming more popular — mobile, is even more popular than the web. Yet, PETs and the companies behind them, do not curate the efficacy of PETs to these platforms. This is mainly due to technical limitations. For instance, blocking advertising and tracking on these platforms are relegated to DNS-based blocking, which has a higher chance of causing breakage for end-users. I will extend the framework of filter rule generation for the web and extend it to create filter rules for other platforms while creating new methodologies to address the limitations of DNS-based blocking.

2. Human-AI Interaction for PETs

End-user Control of PETs. To this day, end-users commonly install PETs that are created and maintained by companies that can have business agendas that may not be in the best interest for the end-user. For instance, adblockers may whitelist ads from companies that have paid to be on that whitelist. My research will establish approaches for end-users to protect their privacy without depending on companies. One possible path would be to design a framework where regular users can contribute to the efficacy of PETs by providing their data in a privacy-preserving and automated manner, while researchers and end-users work jointly to maintain the open-sourced PETs using AI-based approaches.

Auditability of AI-based Tools by Human Experts. Researchers have built AI-based tools to automate tasks such as detecting and blocking advertising and tracking, which reduces the human effort necessary to maintain PETs. However, they have not been widely adopted by industry due to mistrust of AI-based tools in terms of performance (*e.g.*, will they cause breakage) and maintainability (*e.g.*, how to update them when they are not effective). My research will build frameworks and automated tools coupled with user-friendly companion interfaces that allow humans to understand and audit decisions made by AI. This work aims to build trust between humans and AI-based tools and push for industry adoption.

3. Dark Patterns

Scaling Detection and Transparency of Dark Patterns. Dark patterns are deceptive user interfaces within online platforms that can manipulate users into making choices that are not in their best interest. For example, it can manipulate users into agreeing with data collection from advertisers and publishers by providing only an “Accept All” button on a consent popup. Researchers have shown that dark patterns are prevalent across desktop and mobile modalities. However, researchers employ manual efforts to detect dark patterns. I will create methodologies to automatically detect existing and new dark patterns to scale this line of research. In addition, to improve transparency of dark patterns for end-users, my research will develop PETs that highlight and inform users of dark patterns in real-time.

End-user Control of Dark Patterns. Current PETs that deal with dark patterns often make blunt choices for end-users. For instance, a browser extension can automatically close all consent popups for the user upon visiting a website to reduce annoyances. However, this action prevents users from being informed about data collection practices, negating the purpose of consent popups for the public. I plan to explore approaches that can rewrite or remove dark patterns to display neutral user interfaces, so that users can make informed choices for themselves without manipulative interfaces.

References

- [1] Acceptable Ads. Sustainable and non-intrusive advertising. <https://acceptableads.com/>. (Accessed on 01/27/2022).
- [2] eyeo. Ad-Filtering Dev Summit. <https://adfilteringdevsummit.com/>. (Accessed on 09/01/2022).
- [3] eyeo GmbH. Where ad filtering meets profitability. <https://eyeo.com/>. (Accessed on 10/26/2022).
- [4] Hieu Le, Salma Elmalaki, Athina Markopoulou, and Zubair Shafiq. AutoFR: Automated Filter Rule Generation for Adblocking. In *32nd USENIX Security Symposium (USENIX Security 23)*, Anaheim, CA, August 2023. USENIX Association.
- [5] Hieu Le, Athina Markopoulou, and Zubair Shafiq. CV-Inspector: Towards Automating Detection of Adblock Circumvention. In *The Network and Distributed System Security Symposium (NDSS)*, 2021.
- [6] PrivacyCon. OVRSeen Presentation 2022. <https://www.ftc.gov/media/privacycon-2022-part-2>. (Accessed on 10/26/2022).
- [7] PrivacyCon. Smart TV Presentation 2021. <https://www.ftc.gov/media/73491>. (Accessed on 10/26/2022).
- [8] ProperData. Privacy and IoT Research Exploration Workshop 2. <https://properdata.eng.uci.edu/events/privacyiot-workshop-2/>. (Accessed on 10/25/2022).
- [9] Rahmadi Trimananda, Hieu Le, Hao Cui, Janice Tran Ho, Anastasia Shuba, and Athina Markopoulou. OVRseen: Auditing Network Traffic and Privacy Policies in Oculus VR. In *31st USENIX Security Symposium (USENIX Security 22)*, Boston, MA, August 2022. USENIX Association.
- [10] Janus Varmarken, Hieu Le, Anastasia Shuba, Athina Markopoulou, and Zubair Shafiq. The TV is Smart and Full of Trackers: Measuring Smart TV Advertising and Tracking. *Proceedings on Privacy Enhancing Technologies*, 2020(2):129–154, 2020.