I'm not a bot

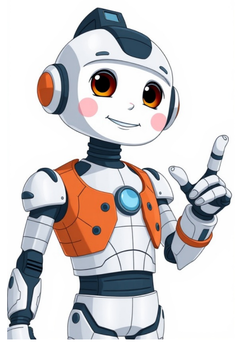

Published by De Gruyter Open / Sciendo before 2022; self-published since Volume 2025 Volume 2024 Volume 2023 Volume 2022 Volume 2021 Volume 2020 Volume 2019 Volume 2018 Volume 2017 Volume 2016 Volume 2015 Privacy Enhancing Technologies Symposium PETS 2014 PETS 2013 PETS 2012 PETS 2011 PETS 2010 PETS 2009 PETS 2008 PETS 2007 Workshop on Privacy Enhancing Technologies PET 2006 PET 2005 PET 2004 PET 2003 PET 2002 Designing Privacy Enhancing Technologies: International Workshop on Design Issues in Anonymity and Unobservability PET 2000 Authors: Andrea Gallardo (Carnegie Mellon University), Chris Choy (Carnegie Mellon University), Jaideep Juneja (Carnegie Mellon University), Efe Bozkir (University of Tübingen), Camille Cobb (University of Illinois), Lujo Bauer (Carnegie Mellon University), Lorrie Cranor (Carnegie Mellon University) Volume: 2023 Issue: 4 Pages: 416–435 DOI: Download PDF Abstract: As technology companies develop mass market augmented reality (AR) glasses that are increasingly sensor-laden and affordable, uses of such devices pose potential privacy and security problems. Though prior work has broadly addressed some of these problems, our work specifically addresses the potential data collection of 15 data types by AR glasses and five potential data uses. Via semi-structured interviews, we explored the attitudes and concerns of 21 current AR technology users regarding potential data collection and data use by hypothetical consumer-grade AR glasses. Participants expressed diverse concerns and suggested potential limits to AR data collection and use, evoking privacy concepts and informational norms. We discuss how participants' attitudes and reservations about data collection and use, like definitions of privacy, are varying and context-dependent, and make recommendations for designers and policy makers, including customizable and multidimensional privacy solutions. Keywords: datasets, neural networks, gaze detection, text tagging Copyright in PoPETs articles are held by their authors. This article is published under a Creative Commons Attribution 4.0 license. Authors: Enze Liu (UC San Diego), Sumanth Rao (UC San Diego), Sam Havron (Cornell Tech), Grant Ho (UC San Diego), Stefan Savage (UC San Diego), Geoffrey M. Voelker (UC San Diego), Damon McCoy (New York University) Volume: 2023 Issue: 1 Pages: 207–224 DOI: Download PDF Abstract: Consumer mobile spyware apps covertly monitor a user's activities (i.e., text messages, phone calls, e-mail, location, etc.) and transmit that information over the Internet to support remote surveillance. Unlike conceptually similar apps used for state espionage, so-called "stalkerware" apps are mass-marketed to consumers on a retail basis and expose a far broader range of victims to invasive monitoring. Today the market for such apps is large enough to support dozens of competitors, with individual vendors reportedly monitoring hundreds of thousands of phones. However, while the research community is well aware of the existence of such apps, our understanding of the mechanisms they use to operate remains ad hoc. In this work, we perform an in-depth technical analysis of 14 distinct leading mobile spyware apps targeting Android phones. We document the range of mechanisms used to monitor user activity of various kinds (e.g., photos, text messages, live microphone access) — primarily through the creative abuse of Android APIs. We also discover previously undocumented methods these apps use to hide from detection and to achieve persistence. Additionally, we document the measures taken by each app to protect the privacy of the sensitive data they collect, identifying a range of failings on the part of spyware vendors (including privacy-sensitive data sent in the clear or stored in the cloud with little or no protection). Keywords: Android Spyware, Android Security, Consumer Spyware Apps, Reverse Engineering, Android API Abuse Copyright in PoPETs articles are held by their authors. This article is published under a Creative Commons Attribution 4.0 license. Authors: Gwendal Patat (Univ Rennes, CNRS, IRISA), Mohamed Sabt (Univ Rennes, CNRS, IRISA), Pierre-Alain Fouque (Univ Rennes, CNRS, IRISA) Volume: 2023 Issue: 4 Pages: 306–321 DOI: Download PDF Abstract: Thanks to HTML5, users can now view videos on Web browsers without installing plug-ins or relying on specific devices. In 2017, W3C published Encrypted Media Extensions (EME) as the first official Web standard for Digital Rights Management (DRM), with the overarching goal of allowing seamless integration of DRM systems on browsers. EME has prompted numerous voices of dissent with respect to the inadequate protection of users. Of particular interest, privacy concerns were articulated, especially that DRM systems inherently require uniquely identifying information on users' devices to control content distribution better. Despite this anecdotal evidence, we lack a comprehensive overview of how browsers have supported EME in practice and what privacy implications are caused by their implementations. In this paper, we fill this gap by investigating privacy leakage caused by EME relying on proprietary and closed-source DRM systems. We focus on Google Widevine because of its versatility and wide adoption. We conduct empirical experiments to show that browsers diverge when complying EME privacy guidelines, which might undermine users' privacy. For instance, we find that many browsers gladly give away the identifying Widevine Client ID with no or little explicit consent from users. Moreover, we characterize the privacy risks of users tracking when browsers miss applying EME guidelines regarding privacy. Because of being closed-source, our work involves reverse engineering to dissect the contents of EME messages as instantiated by Widevine. Finally, we implement EME Track, a tool that automatically exploits bad Widevine-based implementations to break privacy. Keywords: Web Privacy, Web Tracking, DRM, EME, Widevine Copyright in PoPETs articles are held by their authors. This article is published under a Creative Commons Attribution 4.0 license. Authors: Ahmed Roushdy Elkordy (University of Southern California), Jiang Zhang (University of Southern California), Yahya H. Ezzeldin (University of Southern California), Konstantinos Psounis (University of Southern California), Salman Avestimehr (University of Southern California) Volume: 2023 Issue: 1 Pages: 510–526 DOI: Download PDF Abstract: Federated learning (FL) has attracted growing interest for enabling privacy-preserving machine learning on data stored at multiple users while avoiding moving the data off-device. However, while data never leaves users' devices, privacy still cannot be guaranteed since significant computations on users' training data are shared in the form of trained local models. These local models have recently been shown to pose a substantial privacy threat through different privacy attacks such as model inversion attacks. As a remedy, Secure Aggregation (SA) has been developed as a framework to preserve privacy in FL, by guaranteeing the server can only learn the global aggregated model update but not the individual model updates.While SA ensures no additional information is leaked about the individual model update beyond the aggregated model update, there are no formal guarantees on how much privacy FL with SA can actually offer; as information about the individual dataset can still potentially leak through the aggregated model computed at the server. In this work, we perform a first analysis of the formal privacy guarantees for FL with SA. Specifically, we use Mutual Information (MI) as a quantification metric and derive upper bounds on how much information about each user's dataset can leak through the aggregated model update. When using the FedSGD aggregation algorithm, our theoretical bounds show that the amount of privacy leakage reduces linearly with the number of users participating in FL with SA. To validate our theoretical bounds, we use an MI Neural Estimator to empirically evaluate the privacy leakage under different FL setups on both the MNIST and CIFAR10 datasets. Our experiments verify our theoretical bounds for FedSGD, which show a reduction in privacy leakage as the number of users and local batch size grow, and an increase in privacy leakage as the number of training rounds increases. We also observe similar dependencies for the FedAvg and FedProx protocol. Keywords: Federated Learning, Secure Aggregation, Mutual Information, Formal Privacy Guarantee Copyright in PoPETs articles are held by their authors. This article is published under a Creative Commons Attribution 4.0 license. Authors: Alexander Veicht (ETH Zurich), Cedric Renggli (University of Zurich), Diogo Barradas (University of Waterloo) Volume: 2023 Issue: 2 Pages: 188–205 DOI: Download PDF Abstract: Website fingerprinting (WF) attacks, usually conducted with the help of a machine learning-based classifier, enable a network eavesdropper to pinpoint which website a user is accessing through the inspection of traffic patterns. These attacks have been shown to succeed even when users browse the Internet through encrypted tunnels, e.g., through Tor or VPNs. To assess the security of new defenses against WF attacks, recent works have proposed feature-dependent theoretical frameworks that estimate the Bayes error of an adversary's features set or the mutual information leaked by manually-crafted features. Unfortunately, as WF attacks increasingly rely on deep learning and latent feature spaces, our experiments show that security estimations based on simpler (and less informative) manually-crafted features can no longer be trusted to assess the potential success of a WF adversary in defeating such defenses. In this work, we propose DeepSE-WF, a novel WF security estimation framework that leverages specialized kNN-based estimators to produce Bayes error and mutual information estimates from learned latent feature spaces, thus bridging the gap between current WF attacks and security estimation methods. Our evaluation reveals that DeepSE-WF produces tighter security estimates than previous frameworks, reducing the required computational resources to output security estimations by one order of magnitude. Keywords: bayes error, deep neural networks, mutual information, security estimation, traffic analysis, website fingerprinting Copyright in PoPETs articles are held by their authors. This article is published under a Creative Commons Attribution 4.0 license. Authors: Pranav Shriram A (JP Morgan Chase), Nishat Koti (Indian Institute of Science), Varsha Bhat Kukkala (Indian Institute of Science), Arpita Patra (Indian Institute of Science), Bhavish Raj Gopal (Indian Institute of Science), Somya Sangal (Indian Institute of Science) Volume: 2023 Issue: 3 Pages: 24–42 DOI: Download PDF Abstract: Secure shuffle is an important primitive that finds use in several applications such as secure electronic voting, oblivious RAMs, secure sorting, to name a few. For time-sensitive shuffle-based applications that demand a fast response time, it is essential to design a fast and efficient shuffle protocol. In this work, we design secure and fast shuffle protocols relying on the techniques of secure multiparty computation. We make several design choices that aid in achieving highly efficient protocols. Specifically, we consider malicious 3-party computation setting with an honest majority and design robust ring-based protocols. Our shuffle protocols provide a fast online (i.e., input-dependent) phase compared to the state-of-the-art for the considered setting. To showcase the efficiency improvements brought in by our shuffle protocols, we consider two distinct applications of anonymous broadcast and secure graph computation via the GraphSC paradigm. In both cases, multiple shuffle invocations are required. Hence, going beyond standalone shuffle invocation, we identify two distinct scenarios of multiple invocations and provide customised protocols for the same. Further, we showcase that our customized protocols not only provide a fast response time, but also provide improved overall run time for multiple shuffle invocations. With respect to the applications, we not only improve in terms of efficiency, but also work towards providing improved security guarantees, thereby outperforming the respective state-of-the-art works. We benchmark our shuffle protocols and the considered applications to analyze the efficiency improvements with respect to various parameters. Keywords: secure shuffle, anonymous broadcast, secure graph computation, secure multiparty computation Copyright in PoPETs articles are held by their authors. This article is published under a Creative Commons Attribution 4.0 license. Authors: Antreas Dionysiou (University of Cyprus), Vassilis Vassiliades (CYENS Centre of Excellence), Elias Athanasopoulos (University of Cyprus) Volume: 2023 Issue: 1 Pages: 190–206 DOI: Download PDF Abstract: Model Inversion (MI) attacks, that aim to recover semantically meaningful reconstructions for each target class, have been extensively studied and demonstrated to be successful in the white-box setting. On the other hand, black-box MI attacks demonstrate low performance in terms of both effectiveness, i.e., reconstructing samples which are identifiable as their ground-truth, and efficiency, i.e., time or queries required for completing the attack process. Whether or not effective and efficient black-box MI attacks can be conducted on complex targets, such as Convolutional Neural Networks (CNNs), currently remains unclear. In this paper, we present a feasibility study in regards to the effectiveness and efficiency of MI attacks in the black-box setting. In this context, we introduce Deep-BMI (Deep Black-box Model Inversion), a framework that supports various black-box optimizers for conducting MI attacks on deep CNNs used for image recognition. Deep-BMI's most efficient optimizer is based on an adaptive hill climbing algorithm, whereas its most effective optimizer is based on an evolutionary algorithm capable of performing an all-class attack and returning a diversity of images in a single run. For assessing the severity of this threat, we utilize all three evaluation approaches found in the literature. In particular, we (a) conduct a user study with human participants, (b) demonstrate our actual reconstructions along with their ground-truth, and (c) use relevant quantitative metrics. Surprisingly, our results suggest that black-box MI attacks, and for complex models, are comparable, in some cases, to those reported so far in the white-box setting. Keywords: model inversion, inference attack, security, privacy Copyright in PoPETs articles are held by their authors. This article is published under a Creative Commons Attribution 4.0 license. Authors: Akshaye Shenoi (National University of Singapore), Prasanna Karthik Vairam (National University of Singapore), Kanav Sabharwal (National University of Singapore), Jialin Li (National University of Singapore), Dinil Mon Divakaran (Acronis Research) Volume: 2023 Issue: 2 Pages: 206–220 DOI: Download PDF Abstract: IoT devices constantly communicate with servers over the Internet, allowing an attacker to extract sensitive information by passively monitoring the network traffic. Recent research works have shown that a network attacker with a trained machine learning (ML) model can accurately fingerprint IoT devices learned from the (encrypted) traffic flows of the devices. Such fingerprinting attacks are capable of revealing the make and model of the devices, which can further be used to extract detailed user activities. In this work, we develop and propose iPET, a novel adversarial perturbation-based traffic modification system that defends against fingerprinting attacks. iPET design employs GAN (Generative Adversarial Networks) in a tuneable way, allowing users to specify the maximum bandwidth overhead they are willing to tolerate for the defense. A fundamental idea of iPET is to deliberately introduce stochasticity between model instances. This approach limits a counter attack, as it inhibits an attacker from recreating an identical perturbation model and using it for fingerprinting. We evaluate the effectiveness of our defense against state-of-the-art fingerprinting models and with three different attacker capabilities. Our evaluations on synthetic and real-world datasets demonstrate that iPET decreases the accuracy of even the potent attackers. We also show that the traffic perturbations generated by iPET generalize well to different fingerprinting schemes that an attacker may deploy. Keywords: IoT, privacy, fingerprinting, deep learning, adversarial machine learning Copyright in PoPETs articles are held by their authors. This article is published under a Creative Commons Attribution 4.0 license.