Mobile Operandi: Behavioral-Based Fingerprinting Using Permissionless HTML5 APIs

Nathan Reitinger and Michelle L. Mazurek
University of Maryland

Abstract—**Smartphone sensor data, often unguarded by browser permissions, can be used for device fingerprinting. We hypothesize that it is not just hardware, but user behavior, that can be fingerprinted. If possible, this introduces a novel privacy threat, one which is creepily persistent, enables people-based rather than device-based tracking, and allows for cross-device or multi-user, same-device tracking. We propose to collect and analyze sensor data to understand when and how such behavioral fingerprinting may be possible.**

In 2018, almost 4% of the top 100K websites listed on Alexa Top Sites were observed capturing raw, sensor-based data generated from a smartphone's Internal Measurement Unit (IMU) [1]. Gyroscope, accelerometer, ambient light, and proximity sensors were all accessed in the wild, easily available via simple APIs outlined by the HTML5 specification. Somewhat surprisingly, many browsers do not restrict access to these sensors. This opens the door for the stateless identification of website visitors, as is suggested by Das et al. [1], by measuring sensor imperfections introduced in the manufacturing process [2]–[6]. We posit that it is not just the hardware which has identifiable, unique, and repetitive qualities, but users themselves—a new privacy threat not encompassed by current taxonomies of mobile sensor-based attacks [7].

Research Questions.

1) Can permissionless sensors be used for behavioral fingerprinting?
2) Which sensors are most effective for this purpose?
3) What activities are most amenable to such tracking (e.g., reading news, shopping, viewing pictures)?
4) What technical defenses would ameliorate this threat?

Proposed Approach.

HTML5 Permissions. We first characterized permissions needed to access IMU data for different hardware and software configurations. To do this, we built a website housing JavaScript which pulls down touch (`touchstart`, `touchmove`, and `touchend`), motion (`devicemotion`), and orientation (`deviceorientation`) data. We tested this website from varying devices, finding that touch-based data lacks permissions entirely, while orientation and motion data is available for Chrome users, but not Tor or Safari users.[1]

Feature Mining. We will next collect raw IMU data (i.e., touch, motion, and orientation) and use existing human activity recognition datasets to identify feature vectors to be used in classification. For example, touch data will take the form of one 30-length feature vector per "swipe" with attributes including start and stop coordinates, distance, velocity, acceleration, and directional information. These features have been shown to work well in authentication schemes [8]. Notably, although our work has similarities with sensor-based authentication schemes, the limitations of the online environment make these schemes impractical for an off-the-shelf application. In an online environment, a learning phase cannot be used to establish ground truth, and users cannot be relied on to perform predetermined, otherwise uncommon, high-entropy gestures [9]–[12]. Instead, we look to unsupervised techniques for classification of a user's unmediated behavior, collecting IMU events as a surreptitious background process.

Algorithm Assessment. For classification, we will focus on being able to accurately identify repeat visitors; more specifically, whether an incoming data stream (i.e., a user who is currently browsing the website) has characteristics that match existing data associated with a "seen-before" user. In this way, the system could be used for real-time identification of website visitors. We will use a variety of algorithms to assess the accuracy of the match, focusing on unsupervised methods given the limitations stated above.

User Study. We will also differentiate what type of data is most effective (e.g., whether touch data alone can be used for behavioral fingerprinting or, as we hypothesize, if touch and motion data must be combined in order to reach satisfactory accuracy levels) and whether user activity plays a role in identification. Our intuition is that certain types of browsing lead to more accurate identification of users given varying degrees of interaction (e.g., shopping versus reading). To ground the performance of our classifiers and test these hypotheses, we will conduct a longitudinal user study.

Defenses. Finally, we will assess technical solutions for blocking or weakening behavioral fingerprinting. Typical techniques from the tracking-blocking literature, such as spoofing, are likely well suited for lowering classifier accuracy, but must be carefully designed to successfully thwart classification without degrading user experience (e.g., emulating random orientation events may trick a classifier, but should not cause the user to experience unexpected orientation shifts) [7], [13].

Anticipated Outcomes. This work will serve as a proof of concept for the viability of this new privacy threat to mobile browsing. Based on the results, we will recommend that sensors be guarded by permissions in Chrome and that browsers follow Tor and Safari's lead regarding restricted access, rather than permitting sensor access by default.

[1] Apple (Chrome, Safari, and Firefox) requires permissions for orientation data on iOS13+. Android (Chrome and Firefox), tested on v9 (Pie) and v7 (Nougat), allows permissionless access. Android (Tor) restricts permission.

Acknowledgements. This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

References

[1] A. Das, G. Acar, N. Borisov, and A. Pradeep, "The web's sixth sense: A study of scripts accessing smartphone sensors," in *CCS*, 2018.

[2] A. Das, N. Borisov, E. Chou, and M. H. Mughees, "Smartphone fingerprinting via motion sensors: Analyzing feasibility at large-scale and studying real usage patterns," *ArXiv*, June 2016. [Online]. Available: https://arxiv.org/pdf/1605.08763.pdf

[3] A. Das, N. Borisov, and M. Caesar, "Tracking mobile web users through motion sensors: Attacks and defenses," in *NDSS*, 2016.

[4] E. Miluzzo, A. Varshavsky, S. Balakrishnan, and R. R. Choudhury, "Tapprints: Your finger taps have fingerprints," in *MobiSys*, 2012.

[5] S. Dey, N. Roy, W. Xu, R. R. Choudhury, and S. Nelakuditi, "Accelprint: Imperfections of accelerometers make smartphones trackable," in *NDSS*, 2014.

[6] T. Van Goethem, W. Scheepers, D. Preuveneers, and W. Joosen, "Accelerometer-based device fingerprinting for multi-factor mobile authentication," in *International Symposium on Engineering Secure Software and Systems*, 2016.

[7] M. Diamantaris, F. Marcantoni, S. Ioannidis, and J. Polakis, "The seven deadly sins of the HTML5 WebAPI: A large-scale study on the risks of mobile sensor-based attacks," *ACM Transactions on Privacy and Security*, vol. 23, no. 4, Jul. 2020.

[8] M. Frank, R. Biedert, E. Ma, I. Martinovic, and D. Song, "Touchalytics: On the applicability of touchscreen input as a behavioral biometric for continuous authentication," *IEEE Transactions on Information Forensics and Security*, vol. 8, no. 1, pp. 136–148, 2012.

[9] H. Wang, D. Lymberopoulos, and J. Liu, "Sensor-based user authentication," in *European Conference on Wireless Sensor Networks*, 2015.

[10] W.-H. Lee and R. B. Lee, "Implicit smartphone user authentication with sensors and contextual machine learning," in *NDSS*, 2017.

[11] S. Amini, V. Noroozi, A. Pande, S. Gupte, P. S. Yu, and C. Kanich, "Deepauth: A framework for continuous user re-authentication in mobile apps," in *CIKM*, 2018.

[12] R. Masood, B. Z. H. Zhao, H. J. Asghar, and M. A. Kaafar, "Touch and you're trapp(ck)ed: Quantifying the uniqueness of touch gestures for tracking," *Proceedings on Privacy Enhancing Technologies*, vol. 2, pp. 122–142, 2018.

[13] P. Laperdrix, N. Bielova, B. Baudry, and G. Avoine, "Browser fingerprinting: A survey," *ACM Transactions on the Web*, vol. 14, no. 2, pp. 1–33, 2020.