

Stealing PINs via Mobile Sensors: Actual Risk versus User Perception

Maryam Mehrnezhad, Ehsan Toreini, Siamak F. Shahandashti, Feng Hao
School of Computing Science, Newcastle University, UK
Email: {m.mehrnezhad, ehsan.toreini, siamak.shahandashti, feng.hao}@ncl.ac.uk

Abstract—In the first part of this paper, we propose PINlogger.js which is a JavaScript-based side channel attack revealing user PINs on an Android mobile phone. In this attack, once the user visits a website controlled by an attacker, the JavaScript code embedded in the web page starts listening to the motion and orientation sensor streams without needing any permission from the user. By analysing these streams, it infers the user’s PIN using an artificial neural network. Based on a test set of fifty 4-digit PINs, PINlogger.js is able to correctly identify PINs in the first attempt with a success rate of 82.96%, which increases to 96.23% and 99.48% in the second and third attempts respectively. The high success rates of stealing user PINs on mobile devices via JavaScript indicate a serious threat to user security.

In the second part of the paper, we study users’ perception of the risks associated with mobile phone sensors. We design user studies to measure the general familiarity with different sensors and their functionality, and to investigate how concerned users are about their PIN being discovered by an app that has access to all these sensors. Our results show that there is significant disparity between the actual and perceived levels of threat with regard to the compromise of the user PIN. We discuss how this observation, along with other factors, renders many academic and industry solutions ineffective in preventing such side channel attacks.

I. INTRODUCTION

Smartphones equipped with many different sensors such as GPS, light, orientation and motion are continuously providing more features to end users in order to interact with their real-world surroundings. Developers can have access to the mobile sensors either by 1) writing native code using mobile OS APIs [16], 2) recompiling HTML5 code into a native app [32], or 3) using standard APIs provided by the W3C which are accessible through JavaScript code within a mobile browser¹. The last method has the advantage of not needing any app-store approval for releasing the app or doing future updates. More importantly, the JavaScript code is platform independent, i.e., once the code is developed it can be executed within any modern browser on any mobile OS.

In-browser access risks. While sensor-enabled mobile web applications provide users more functionalities, they raise

¹w3.org/TR/#tr_Javascript_APIs

Permission to freely reproduce all or part of this paper for noncommercial purposes is granted provided that copies bear this notice and the full citation on the first page. Reproduction for commercial purposes is strictly prohibited without the prior written consent of the Internet Society, the first-named author (for reproduction of an entire paper only), and the author’s employer if the paper was prepared within the scope of employment.
EuroUSEC ’16, 18 July 2016, Darmstadt, Germany
Copyright 2016 Internet Society, ISBN 1-891562-45-2
http://dx.doi.org/10.14722/eurosec.2016.23008

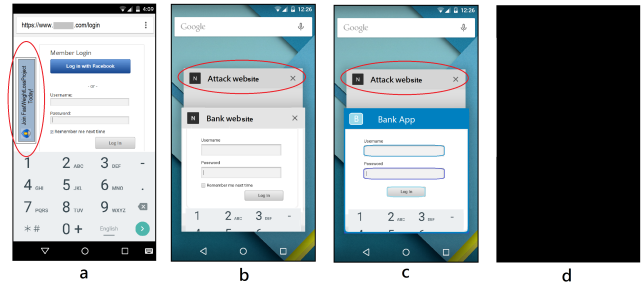


Fig. 1. PINlogger.js potential attack scenarios; a) the malicious code is loaded in an iframe and the user is on the same tab, b) the attack tab is already open and the user is on a different tab, c) the attack content is already open in a minimised browser, and the user is on an installed app, d) the attack content is already open in a (minimised) browser, and the screen is locked. The attacker listens to the side channel motion and orientation measurements of the victim’s mobile device through JavaScript code, and uses machine learning methods to discover the user’s sensitive information such as activity types and PINs.

new privacy and security concerns. Both the academic community and the industry have recognised such issues regarding certain sensors such as geolocation [18]. For a website to access the geolocation data, it must ask for explicit user permission. However, to the best of our knowledge, there is little work evaluating the risks of in-browser access to other sensors. Unlike in-app attacks, an in-browser attack, i.e., via JavaScript code embedded in a web page, does not require any app installation. Furthermore, JavaScript code does not require any user permission to access sensor data such as device motion and orientation. Furthermore, there is no notification while JavaScript is reading the sensor data stream. Hence, such in-browser attacks can be carried out far more covertly than the in-app counterparts. However, an effective in-browser attack still has to overcome the technical challenge that the sampling rates available in browser are much lower than those in app. For example, as we observed in [22], frequency rates of motion and orientation sensor data available in-browser are 3 to 5 times lower than those of accelerometer and gyroscope available in-app.

Motion and orientation sensors detail. According to W3C specifications [1] motion and orientation sensor data are a series of different measurements:

- device *orientation* which provides the physical orientation of the device, expressed as three rotation angles (α , β , γ) in the device’s local coordinate frame,
- device *acceleration* which provides the physical acceleration of the device, expressed in Cartesian coordinates (x , y , z) in the device’s local coordinate frame,

- device *acceleration-including-gravity* which is similar to acceleration except that it includes gravity as well
- device *rotation rate* which provides the rotation rate of the device about the local coordinate frame, expressed as three rotation angles (α , β , γ), and
- *interval* which provides the constant sampling rate and is expressed in milliseconds (ms).

The device coordinate frame is defined with respect to the standard position of the mobile screen. When it is in the portrait mode, x and y axes are in the plane of the screen and are positive towards the screen’s right and up, and z is perpendicular to the plane of the screen and is positive outwards from the screen. Moreover, the sensor data discussed above are processed sensor data obtained from multiple physical sensors such as gyroscope and accelerometer. In the rest of this paper, unless specified otherwise, by sensor data we mean the sensor data accessible through mobile browsers which includes acceleration, acceleration-including-gravity, rotation rate, and orientation.

Motivation. Many popular browsers such as Safari, Chrome, Firefox, Opera and Dolphin have already implemented access to the above sensor data. As we demonstrated in [21] and [22], all of these mobile browsers allow such access when the code is placed in any part of the active tab including *iframes* (Figure 1, a). In some cases such as Chrome and Dolphin on iOS, an inactive tab including the sensor listeners have access to the sensor measurements as well (Figure 1, b). Even worse, some browsers such as Safari allow the inactive tabs to access the sensor data, when the browser is minimised (Figure 1, c), or even when the screen is locked (Figure 1, d). Mobile operating systems and browsers do not seem to be implementing consistent access control policies in regard to mobile orientation and motion sensor data. Furthermore, W3C specifications [1] do not discuss any risks associated with this potential vulnerability. Because of the low sampling rates available in browser, the community have been neglecting the security risks associated with in-browser access to such sensor data. However, in TouchSignatures [22], we showed that despite the low sampling rates, it is possible to identify user touch actions such as click, scroll, and zoom and even the numpad’s digits. In this work we contribute to the study of such attacks as follows:

- We introduce PINLogger.js, an attack on full 4-digit PINs as opposed to only single digits in [22]. We show that unregulated access to these sensors impose more serious security risks to the users in comparison with more well-known sensors such as camera, light and microphone.
- We conduct user studies to investigate users’ understanding about these sensors and also their perception of the security risks associated with them. We show that users in fact have fewer security concerns about these sensors comparing to more well-known ones.
- We study and challenge current suggested solutions, and discuss why our studies show they cannot be effective. We argue that a usable and secure solution is not straightforward and requires further research.

II. PINLOGGER.JS

In this section, we describe an advanced attack on user’s PINs by introducing PINlogger.js. In the following subsections, we describe the attack approach, our program implementation, data collection, feature extraction, and neural network.

A. Attack approach

We consider an attacker who wants to learn the user’s PIN tapped on a soft keyboard of a smartphone via side channel information. We consider (digit-only) PINs since they are popular passwords used by users for many purposes such as unlocking phone, SIM PIN, NFC payments, bank cards, other banking services, gaming, and other personalised applications such as healthcare, insurance, etc. Unlike similar works which have to gain the access through an installed app [23], [27], [24], [10], [29], [30], [26], [33], [3], [11], our attack does not require any user permission. Instead, we assume that the user has loaded the malicious web content in the form of an iframe, or another tab while working with the mobile browser as shown in Figure 1. At this point, the attack code has already started listening to the sensor sequences from the user’s interaction with the phone.

In order to uncover when the user enters his PIN, we need to classify his touch actions such as click, scroll, and zoom. We already have shown in TouchSignatures [22] that with the same sensor data and by applying classification algorithms, it is possible to effectively identify user’s touch actions. Here, we consider a scenario after the touch action classification. In other words, our attacker already knows that the user is entering his PIN. Moreover, unless explicitly noted, we consider a generic attack scenario which is not user-dependant. This means that we do not need to train our machine learning algorithm with the same user as the subject of the attack. Instead, we have a one-round training phase with data from multiple voluntary users, and use the obtained trained algorithm to output other users’ PINs later. This approach has the benefit of not needing to trick individual users to collect data for training.

B. Web program implementation

We implemented a web page with embedded JavaScript code in order to collect the data from voluntary users. Our code registers two listeners on the window object to have access to orientation and motion data, separately. The event handlers defined for these purposes are named *DeviceOrientationEvent* and *DeviceMotionEvent*, respectively. On the client side, we developed a GUI in HTML5 which shows random 4-digit PINs to the users and activates a numpad for them to enter the PINs as shown in Figure 2. All sensor sequences are sent to the database along with their associated labels which are the digits of the entered PINs. We implemented our server program using Node.js (nodejs.org). Our code sends the orientation and motion sensor data of the mobile device to our NoSQL database using MongoLab (mongolab.com, web-based service for MongoDB). When the event listener fires, it establishes a socket by using Socket.IO (socket.io) between the client and the server and constantly transmits the sensor data to the database. Both Node.js and MongoDB (as a document-oriented database) are known for being capable of supporting data intensive applications in real time.




Fig. 2. Different input methods used by the users for PIN entrance.

C. Data collection

Following the approach of Aviv et al. [3] and Spreitzer [30], we consider a set of 50 fixed random PINs in this paper. We conducted our user studies using Chrome on an Android device (Nexus 5). The experiments and results are based on the collected data from 5 users, each entering all the 50 4-digit PINs for 5 times. Our voluntary participants were university students and staff and performed the experiments at university offices. We simply explained to them that all they needed was to enter a few PINs shown in a web page.

In relation to the environmental setting for the data collection, we asked the users to remain sitting in a chair while working with the phone. We did not require our users to hold the phone in any particular mode (portrait or landscape) or work with it by using any specific input method (using one or two hands). We let them choose their most comfortable posture for holding the phone and working with it as they do in their usual manner. While watching the users during the experiments, we noticed that all of our users used the phone in the portrait mode by default. Users were either leaning their hands on the desk or freely keeping them in the air. We also observed the following input methods used by the users.

- Holding the phone in one hand and entering the PIN with the thumb of the same hand (Figure 2, left).
- Holding the phone in one hand and entering the PIN with the fingers of the other hand (Figure 2, centre).
- Holding the phone with two hands and entering the PIN with the thumbs of both hands (Figure 2, right).

In the first two cases, users exchangeably used either their right hands or left hands in order to hold the phone. In order to simulate a real world data collection environment, we took the phone to each user’s workspace and briefly explained the experiment to them, and let them complete the experiment without our supervision. All users found this way of data collection very easy and could finish the experiments without any difficulties.

D. Feature extraction

In order to build the feature vector as the input to our classifier algorithm, we consider both time domain and frequency domain features. We improve our suggested feature vectors in [22] by adding some more complex features such as the correlation between the measurements. This addition improves the results, as we will discuss in Section III. As discussed before, 12 different sequences obtained from the collected

data include orientation (ori), acceleration (acc), acceleration-including-gravity (accG), and rotation rate (rotR) with three sequences (either x, y and z, or α , β and γ) for each sensor measurement. As a pre-processing step and in order to remove the effect of the initial position and orientation of the device, we subtract the initial value in each sequence from subsequent values in the sequence.

We use these pre-processed sequences for feature extraction in time domain directly. In frequency domain, we apply the Fast Fourier transform (FFT) on the pre-processed sequences and use the transformed sequences for feature extraction. In order to build our feature vector, first we obtain the maximum, minimum, and average values of each pre-processed and FFT sequences. These statistical measurements give us $3 \times 12 = 36$ features in the time domain, and the same number of features in the frequency domain. We also consider the total energy of each sequence in both time and frequency domains calculated as the sum of the squared sequence values, i.e., $E = \sum v_i^2$ which gives us 24 new features.

The next set of features are in time domain and are based on the correlation between each pair of sequences in different axes. We have 4 different sequences; ori, acc, accG, and rotR, each represented by 3 measurements. Hence, we can calculate 6 different correlation values between the possible pairs; (ori, acc), (ori, accG), (ori, rotR), (acc, accG), (acc, rotR), and (accG, rotR), each presented in a vector with 3 elements. We use the Correlation coefficient function in order to calculate the similarity rate between the mentioned sequences. The correlation coefficient method is commonly used to compare the similarity of the shapes of two signals (e.g. [5]). Given two sequences A and B and $Cov(A, B)$ denoting covariance between A and B , the correlation coefficient is computed as below:

$$R_{AB} = \frac{Cov(A, B)}{\sqrt{Cov(A, A) \cdot Cov(B, B)}} \quad (1)$$

The correlation coefficient of two vectors measures their linear dependence by using covariance. By adding these new 18 features, our feature vector consists of a total of 114 features.

E. Neural network

We apply a supervised machine learning algorithm by using an Artificial Neural Network (ANN) to solve this classification problem. The input of an ANN system could be either raw data, or pre-processed data from the samples. In our case, we have preprocessed our samples by building a feature vector as described before. Therefore, as input, our ANN receives a set of 114 features for each sample. As explained before, we collected 5 sample per each 4-digit PINs from 5 different users, giving us 1250 feature vectors in general.

The feature vectors are mapped to specific labels from a finite set: i.e., 50 fixed random 4-digit PINs. We train and validate our algorithm with two different subsets of our collected data, and test the neural network against a separate subset of the data. We train the network with 70% of our data, validate it with 15% of the records and test it with the remaining 15% of our data set. We use a pattern recognition/classifying network in Matlab with one hidden layer and 1000 nodes.

| Attempts | Identification rate |
|----------|---------------------|
| One | 82.96% |
| Two | 96.23% |
| Three | 99.48% |

TABLE I. PINLOGGER.JS’S PIN IDENTIFICATION RATES IN DIFFERENT ATTEMPTS.

| Attempts | User independent | User dependent |
|----------|------------------|----------------|
| One | 71.57% | 80.21% |
| Two | 82.83% | 90.24% |
| Three | 92.01% | 95.05% |

TABLE II. AVERAGE DIGIT IDENTIFICATION RATES IN DIFFERENT ATTEMPTS.

Pattern recognition/classifying networks normally use a scaled conjugate gradient (SCG) back-propagation algorithm for updating weight and bias values in training. Scaled conjugate gradient is a fast supervised learning algorithm [25].

III. EVALUATION

In this section we present the results of our attack and compare them with other works.

A. PINlogger.js success rate

Table I shows the accuracy of our ANN trained with the data from all users. Since these results are based on the collected data from all users, we refer to it as the user-independent mode. As the table shows, in the first attempt PINlogger.js is able to infer the user’s 4-digit PIN correctly with accuracy of 82.96%. The results get better on further attempts. As the table shows, our system is able to reveal the user’s PIN with nearly 100% accuracy in three attempts. By comparison, a random attack can guess a PIN from a set of 50 PINs with the probability of 2% in the first attempt, and 6% in three attempts.

B. User-dependent mode

In order to study the impact of individual training, we trained, validated and tested the network with the data collected from one user. We refer to this mode of analysis as the user-dependent mode. We asked our user to enter 50 random PINs, each five times, and repeat the experiment for 5 times (rounds). The reason we have repeated the experiments is that the classifier needs to receive enough samples to be able to train the system. Interestingly, our user used all three different input methods shown in Figure 2 during the PIN entrance. As expected, our classifier performs better when it is personalized: the accuracy increases to 91.42% in the first attempt, and 98.64% and 100% in two and three attempts, respectively.

In the user-dependent mode, convincing the users to provide the attacker with sufficient data for training customised classifiers is not easy, but still possible. Approaches similar to gaming apps such as Math Trainer² could be applied. Math-based CAPTCHAs are possible web-based alternatives. Any other web-based game application which segments the GUI similar to a numerical keypad will do as well. Nonetheless, this is out of the scope of this paper since we mainly follow a user-independent approach.

C. Guessing the PIN from the entire PIN space

One might argue that the attack should be evaluated against the whole 4-digit PIN space. However, we believe that the attack could still be practical when selecting from a limited set of PINs since users do not select their PINs randomly [8].

It has been reported that around 27% of all possible 4-digit PINs belong to a set of 20 PINs³, including straightforward ones like ‘1111’, ‘1234’, or ‘2000’. Nevertheless, we present the results of our analysis of the attack against the entire search space for both the user-independent and user-dependent modes.

For user-independent mode, we trained another ANN in order to infer a single digit on the numpad. In this experiment, we considered 10 classes of the entered digits (0–9) from the data we collected on 4-digit PINs used in Section III-A. For user-dependent mode, we trained personalised classifiers for each user. Unlike the test condition of Section III-B, we did not have to increase the number of rounds of PIN entry here since we had enough samples for each digit per user. Hence in the user-dependent mode in this section, we used the average of the results of our 5 users. The average identification rates of different digits are presented in Table II.

The results in our user-independent mode show that it is possible to correctly infer digits in over 71% of the cases in the first attempt, going up to 92% in three attempts. This means that for a 4-digit PIN and based on the obtained sensor data, the attacker can guess the PIN to be within a set of $3^4 = 81$ possible PINs with a probability of success equal to $0.92^4 = 71.67\%$. A random attack, however, can only predict the 4-digit PIN with the probability of 0.81% in 81 attempts. By comparison, PINlogger.js achieves a dramatically higher success rate than a random attacker. Using a similar argument, in the user-dependent mode the success probability of guessing the PIN in 81 attempts is 81.62%. In the same setting, Cai and Chen report a success rate of 65% using accelerometer and gyroscope data [2] and Simon and Anderson’s PIN Skimmer only achieves a 12% success rate in 81 attempts using camera and microphone [29]. Our results in digit recognition in this paper are also better than what is achieved in TouchSignatures [22]. In summary, PINlogger.js performs better than all sensor-based digit-identifier attacks in the literature.

D. Comparison with related work

Obtaining sensitive information about users such as PINs based on mobile sensors through a malicious app running in the background has been actively explored by researchers in the field. For example, GyroPhone, by Michalevsky et al. [23], shows that gyroscope data is sufficient to identify the speaker and even parse speech to some extent. Other examples include Accessory [27] by Owusu et al. and Tappprints [24] by Miluzzo. They infer passwords on full alphabetical soft keyboards based on accelerometer measurements. Touchlogger [10] is another example by Cai and Chen [2] which shows the possibility of distinguishing user’s input on a mobile numpad by using accelerometer and gyroscope. The same authors demonstrate a similar attack in [11] on both numerical and full keyboards. The only work which relies on in-browser access to sensors to

²play.google.com/store/apps/details?id=com.solirify.mathgame

³datagenetics.com/blog/september32012/

| Featured Work | PIN Skimming [30] | PIN Skimmer [29] | Keylogging by Mic [26] | TapLogger [33] | Acc. side channel [3] | PINlogger.js | |
|---------------------|-------------------|--------------------|------------------------|-----------------|-----------------------|------------------------|----------------|
| Sensor Access type | Light in-app | Camera, Mic in-app | Mic, Gyr in-app | Acc, Ori in-app | Acc in-app | Motion, Ori in-browser | |
| Training approach | user-dependent | user-dependent | user-dependent | user-dependent | user-independent | user-independent | user-dependent |
| Identification rate | | | | | | | |
| First attempt | NA | NA | 94% | 40% | 18% | 82.96% | 91.42% |
| Second attempt | 50% | 30% | NA | 75% | NA | 96.23% | 98.64% |
| Fifth attempt | 65% | 50% | NA | 100% | 43% | 100% | 100% |

TABLE III. COMPARISON OF PINLOGGER.JS WITH RELATED ATTACKS ON 4-DIGIT PINs.

attack a numpad is our previous work, TouchSignatures [22]. All of these works, however, aim for the individual digits or characters of a keyboard, rather than the entire PIN or password.

Another category of works directly target user PINs. For example, PIN skimmer by Simon and Anderson [29] is an attack on a user’s numpad and PINs using the camera and microphone on the smartphone. Spreitzer suggests another PIN Skimming attack [30] and steals a user’s PIN based on the measurements from the smartphone’s ambient light sensor. Narain et al. introduce another attack [26] on smartphone numerical and alphabetical keyboards and the user’s PINs and credit card numbers by using the smartphone microphone. TapLogger by Xu et al. [33] is another attack on the smartphone numpad which outputs the pressed digits and PINs based on accelerometer and orientation sensor data. Similarly, Aviv et al. introduce an accelerometer-based side channel attack on the user’s PINs and patterns in [3]. We choose to compare PINlogger.js with the works in this category since they have the same goal of revealing the user’s PINs. Table III presents the results of our comparison.

As shown in Table III, PINlogger.js is the only attack on PINs which acquires the sensor data via JavaScript code. In-browser JavaScript-based attacks impose even more security threats to users since unlike in-app attacks, they do not require any app installation and user permission to work. Moreover, the attacker does not need to develop different apps for different platforms such as Android, iOS, and Windows. Once the attacker develops the JavaScript code, it can be deployed to attack all mobile devices regardless of the platform. Moreover, Touchlogger.js and [3] are the only user-independent works. By contrast, the results from other works are based on training the classifiers for individual users. In other words, they assume the attacker is able to collect input training data from the victim user before launching the PIN attack. We do not have such an assumption as the training data is obtained from all users in the experiment. In terms of accuracy, with the exception of [26], PINlogger.js generally outperforms other works with an identification rate of 82.96% in the first try, and 96.23% and 100% in the second and fifth attempts, respectively. This is a significant success rate (despite that the sampling rate in-browser is much lower than that available in-app) and confirms that the described attack imposes a serious threat to the users’ security and privacy.

IV. WHY DOES THIS VULNERABILITY EXIST?

Although reports of side channel attacks based on the in-browser access to mobile sensors via JavaScript are relatively recent, similar attacks via in-app access to mobile sensors have been known for years. Yet the problem has not been fixed.

We believe a contributing factor is that users seem to be less familiar with the relatively newer (and less advertised) sensors such as motion and orientation, as opposed to their immediate familiarity with well-established sensors such as camera and GPS. For example, a user has asked this question on a mobile forum: “... What benefits do having a gyroscope, accelerometer, proximity sensor, digital compass, and barometer offer the user? I understand it has to do with the phone orientation but am unclear in their benefits. Any explanation would be great! Thanks!”⁴.

We design and conduct user studies in this work in order to investigate to what extent are these sensors and their risks known to the users.

A. List of mobile sensors

We prepared a list of different mobile sensors by inspecting the official websites of the latest iOS and Android products, and the specifications that W3C and Android provide for developers. We also added some extra sensors as common sensing mobile hardware which are not covered before.

- iPhone 6⁵: Touch ID, Barometer, Three-axis gyro, Accelerometer, Proximity sensor, Ambient light sensor.
- Nexus 6P⁶: Fingerprint sensor, Accelerometer, Gyroscope, Barometer, Proximity sensor, Ambient light sensor, Hall sensor, Android Sensor hub.
- Android [16]: Accelerometer, Ambient temperature, Gravity (software or hardware), Gyroscope, Light, Linear Acceleration (software or hardware), Magnetic Field, Orientation (software), Pressure, proximity, Relative humidity, Rotation vector (Software or Hardware), Temperature.
- W3C⁷ [1]: Device orientation (software), Device motion (software), Ambient light, Proximity, Ambient temperature, Humidity, Atmospheric Pressure.
- Extra sensors (Common sensing hardware): Wireless technologies (WiFi, Bluetooth, NFC), Camera, Microphone, Touch screen, GPS.

Unless specified otherwise, all the listed sensors are hardware sensors. We added the last category of the sensors to this list since they indeed sense the device’s surrounding although in different ways. However, they are neither counted as sensors

⁴forums.androidcentral.com/verizon-galaxy-nexus/171482-barometer-accelerometer-how-they-useful.html

⁵apple.com/uk/iphone-6/specs/

⁶store.google.com/product/nexus_6p

⁷w3.org/2009/dap/

in mobile product descriptions, nor in technical specifications. These sensors are often categorised as OS resources [31], and hence different security policies apply to them.

B. User study

We prepared a list of sensors based on the above. We asked volunteer participants to rate the level of their familiarity with each sensor. In all of our studies, we had 30 participants (13 self-identified as male and 17 as female) recruited from the university and local community through social and vocational networks, from 18 to 59 years old, with a median age of 31. Except one, none of the participants were studying or working in the field of computer security. Our university participants were from multiple degree programs and levels, and the remaining participants worked in a different range of fields. Moreover, our participants owned a wide range of mobile devices, and had been using a smartphone/tablet for 6 years on average (from 0 to 11 years). We interviewed our participants at a university office and gave each an Amazon voucher (worth £10) at the end for their participation. Details of the interview template can be found in the Appendix.

For a list of 25 different sensors, we used a five-point scale self-rated familiarity questionnaire as used in [19]: “I’ve never heard of this”, “I’ve heard of this, but I don’t know what this is”, “I know what this is, but I don’t know how this works”, “I know generally how this works”, and “I know very well how this works”. The list of sensors was randomly ordered for each user to minimize bias. In addition, we needed to observe the experiments to make sure users were answering the questions based on their own knowledge in order to avoid the effect of processed answers. Full descriptions of all studies are provided in the Appendix. Fig. 3 summarizes the results of this study.

Our participants were generally surprised to hear about some sensors and impressed by the variety. As one may expect, newer sensors tend to be less known to the users in comparison to older ones. In particular, our participants were generally not familiar with ambient sensors. Also low-level hardware sensors such as accelerometer and gyroscope, seem to be less known to the users in comparison with high-level software ones such as motion, orientation, and rotation. We suspect that this is partly due to the fact that the high-level sensors are named after their functionalities and can be more immediately related to user activities.

We also noticed that a few of the participants knew some of the low-level sensors by name but they could not link them to their functionality. For example, one of our participants which knew almost all of the listed sensors (except hall sensor and sensor hub) stated that: “When I want to buy a mobile [phone], I do a lot of search, that is why I have heard of all of these sensors. But, I know that I do not use them (like accelerometer and gyroscope)”.

On the other hand, as the functionalities of mobile devices grow, vendors quite naturally turn to promote the software capabilities of their products, instead of introducing the hardware. For example, many mobile devices are recognised for their gesture recognition features by the users, however the same users might not know how these devices provide such a feature. For instance, one of the participants commented on a

feature on her smartphone called “Smart Stay”⁸ as follows: “I have another sensor on my phone: Smart Stay. I know how it works, but I don’t know which sensors it uses”.

V. RISK PERCEPTION OF MOBILE SENSORS

In this section, we study the participants’ risk perception of mobile sensors. There have been several studies on risk perception addressing different aspects of mobile technology. Some works discuss the risks that users perceive on smartphone authentication methods such as PINs and patterns [17], TouchID and Android face unlock [14], and implicit authentication [20]. Other works focus on the privacy risks of certain sensors such as GPS [4]. In [28], Raji et al. show users’ concerns (on disclosure of selected behaviours and contexts) about a specific sensor-enabled device called AutoSense⁹. To the best of our knowledge, the research presented in this paper is the first that studies the user risk perception for a comprehensive list of mobile sensors (25 in total). We limit our study to the level of perceived risks users associate with their PINs being discovered by each sensor. The reasons we chose PINs are that first, finding one’s PIN is a clear and intuitive security risk, and second, we can put the perceived risk levels in context with respect to the actual risk levels for a number of sensors as described in Table III.

A. Methodology

For this study, we interviewed the same group of users from Section IV-B in two phases. In phase one, we gave the same sensor list (randomized for each user). We asked users to rate the level of risk they perceive for each sensor in regards to revealing their PINs. We described a specific scenario in which a game app which has access to all these sensors is open in the background and the user is working on his online banking app, entering a PIN. We used a self-rated questionnaire with five-point scale answers following the same terminology as used in [28]: “Not concerned”, “A little concerned”, “Moderately concerned”, “Concerned”, and “Extremely concerned”. During this phase, we asked the users to rely on the information that they already had about each sensor (see the Appendix for details).

In the second phase, first we provided the participants with a short description of each sensor and let them know that they can ask further questions until they feel confident that they understand the functionality of all sensors. Afterwards, we asked the participants to fill in another copy of the same questionnaire on risk perceptions (details in the Appendix). The results are presented in Fig. 4.

B. Intuitive risk perception

We make the following observations from the results of the experiment.

Touch Screen. Although our participants rated touch screen as one of the most risky sensors in relation to a PIN discovery scenario, still about half of our participants were either moderately concerned, a little concerned, or not

⁸samsung.com/us/support/answer/ANS00035658/234302/SCH-R950TSAUSC

⁹sites.google.com/site/autosenseproject/




Fig. 3. Level of self-declared knowledge about different mobile sensors. Question: “To what extent do you know each sensor on a mobile device?” Sensors are ordered based on the aggregate percentage of participants declaring they know generally or very well how each sensor works. This aggregate percentage is shown on the right hand side.

concerned at all. Through our conversations with the users, we received some interesting comments, e.g., “Why any of these sensors should be dangerous on an app while I have officially installed it from a legal place such as Google Play?”, and “As long as the app with these sensors is in the background, I have no concern at all”. It seems that a more general risk model in relation to mobile devices is affecting the users’ perception in regard to the presented PIN discovery threat. This fact can be a topic of research on its own, and is out of the scope of this paper.

Communicational Sensors. One category of the sensors which users are relatively more concerned about includes WiFi, Bluetooth and NFC. For example one of the participants commented that: “I am not concerned with physical [motion, orientation, accelerometer, etc.]/ environmental [light, pressure, etc.] sensors, but network ones. Hackers might be able to transfer my information and PIN”. This comment is understandable since we asked them to what extent they were concerned about *each sensor* in regard to the PIN discovery.

Identity-related Sensors. Another category which has

been rated more risky than others contains those sensors which can capture something related to the user’s identity i.e. fingerprint, TouchID, GPS, camera, and microphone. Despite that we described a PIN-related scenario, our participants were still concerned about these sensors. This was also pointed out by a few participants through the comments. For example a user stated: “..., however, GPS might reveal the location along with the user input PIN that has a risk to reveal who (and where) that PIN belongs to. Also the fingerprint/TouchID might recognize and record the biometrics with the user’s PIN”. Some of these sensors such as GPS, fingerprint, and TouchID, however, can not cause the disclosure of PINs on their own. Hence, the concern does not entirely match the actual risk. Similar to the discussion on touch screen, we believe that a more general risk model on mobile technology influences the users to perceive risk on specific threats such as the one we presented to them.

Environmental Sensors. The level of concern on ambient sensors (humidity, light, pressure, and temperature) is generally low and stays low after the users are provided with the description of the sensors (see Fig. 4). In many cases, our




Fig. 4. Users’ perceived risk for different mobile sensors, before (top bars) and after (bottom bars) being presented with descriptions of sensors. Question: “To what extent are you concerned about each sensor’s risk to your PIN?”. Sensors are ordered based on the aggregate percentage of participants declaring they are either concerned or extremely concerned about each sensor before seeing the descriptions. This aggregate percentage is the first value presented on the right hand side.

users expressed that they were concerned about these sensors simply because they did not know them: “[now that I know these sensors,] I am quite certain that movement/environmental sensors would not affect the security of personal id/passwords etc.”. In fact, researchers have reported that it is possible to infer the user’s PIN using the ambient light sensor data [30], although, to our knowledge, exploits of other environmental sensors have not been reported in the literature.

Movement Sensors. On the sensors related to the movement and the position of the phone (accelerometer, gyroscope, motion, orientation, and rotation), the users display varying

levels of the risk perceptions. In some cases they are slightly more concerned, but in others they are less concerned once they know the functionality. Some of our users stated that since they did not know these sensors, they were not concerned at all, but others were more concerned when they were faced with new sensors. Overall, knowing, or not knowing these sensors has not affected the perceived risk level significantly, and they were rated generally low in both cases.

Motion and Orientation Sensors. The sensors which we used in our attack, namely orientation, rotation, and motion, have not been generally scored high for their risk in revealing

PINs. Users do not seem to be able to relate the risk of these sensors to the disclosure of their PINs, despite that they seem to have an average general understanding about how they work. On hardware sensors such as accelerometer and gyroscope, the risk perception seems to be even lower. A few comments include: “In my everyday life, I don’t even think about these [movement] sensors and their security. There is nothing on the news about their risk”, and “I have never been thinking about these [movement] sensors and I have not heard about their risk”. On the other hand, some of the participants expressed more concerns for sensors that they were familiar with, as one wrote, “You always hear about privacy stuff for example on Facebook when you put your location or pictures”. Similarly, it seems that having a previous risk model is a factor that might explain the correlation between the user’s knowledge and their perceived risk.

C. General knowledge versus risk perception

Figs. 3 and 4 suggest that there may be a correlation between the relative level of knowledge users have about sensors and the relative level of risk they perceive from them. We limit our attention to users’ knowledge before being presented with sensor descriptions. We confirm our observation of correlation using Spearman’s rank-order correlation measure.

Spearman’s correlation between the comparative knowledge (median: “I know what this is, but I don’t know how this works”, IQR: “I’ve never heard of this” – “I know very well how this works”) and the perceived risk about different sensors (median: “Not concerned”, IQR: “Not concerned” – “A little concerned”) was $r = 0.61$ ($p < 0.05$). This result supports that the more the users know about these sensors, the more concern they express about the risk of the sensors revealing PINs. We acknowledge that other methods of ranking the results, e.g. using median, produce slightly different final rankings. However, given the high confidence level of the above test, we expect the correlation to be supported if other methods of ranking are used.

Assuming that customer demand drives better security designs, the above correlation may explain why sensors that are newer to the market have not been considered as OS resources and consequently have not been subject to similar strict access control policies.

D. Perceived risk vs the actual risk

We are specifically interested in the users’ relative risk perception of sensors in revealing their PINs in comparison to the actual relative risk level of these sensors. We list the results reported in the literature in Table III for the following sensors: light, camera, microphone, gyroscope, motion, and orientation. Fig. 4 shows that users generally have expressed more concern about sensors such as camera and microphone than accelerometer, gyroscope, orientation, and motion. This does not match the actual risk levels since the latter sensors allow PIN recovery with higher accuracy as we have shown in Section III. When asked after filling the questionnaire, most participants could not come up with realistic attack scenarios using camera and microphone. For microphone, some users thought they might say the PIN out loud. For camera, a few of our participants thought face recognition might be used to

recover the PIN, hence they rated camera’s risk to their PINs high. One user thought the camera might capture the reflection of the entered PIN in her glasses.

Among our participants, one mentioned but described doubt about motion, orientation, accelerometer, and gyroscope being able to record the shakes of the mobile phone while entering a PIN after they saw the sensor descriptions: “I feel those positional sensors might be able to reveal something about my activities, for example if I open my banking app or enter my PIN. But it is extremely hard for different users, and when working with different hands and positions”. This participant expressed only “a little concern” about them, stating that: “..., and by little concern, I mean extremely little concern”. One of our participants was completely familiar with these attacks and in fact had read some related papers. This user was “extremely concerned”. Other users who rated these sensors risky in general, said they were generally concerned about different sensors. One commented: “I can not think of any particular situation in which these sensors can steal my PIN, but the hackers can do everything these days.”

VI. POSSIBLE SOLUTIONS

In this section, we discuss the current academic and industrial countermeasures to mitigate sensor-based attacks.

A. Academic approach

Different solutions to address the in-app access attacks have been suggested in the literature: e.g., restricting the sensor to one app, reducing the sampling rate, temporal pause of the sensor on sensitive entries such as keyboard, rearranging keyboard for password entrance, asking for explicit permission from the user, ranking apps based on their similarities to malware, and obfuscating anomalies in sensor data [26], [3], [30], [33], [29], [23], [24], [27], [13], [6]. However, after many years of research on showing the serious security risks of sensors such as accelerometer and gyroscope, none of the major mobile platforms have revised their in-app access policy.

We believe that the risks of unmanaged sensors on mobile phones, specially through JavaScript code, are not known very well yet. More specifically, many OS/app level solutions such as asking for permissions at the installation time, or malware detection approaches would not work in the context of a web attack. In our previous work [22], we suggested to apply the same security policies as those for camera, microphone, and GPS for the motion and orientation sensors. Our suggestion was to set a multi-layer access control system on the OS and browser levels. However, the usability and effectiveness of this solution are arguable. First, asking too many permissions from the user for different sensors might not be usable. Furthermore, for some basic use cases such as gesture recognition to clear a web form, or adjusting the screen from portrait to landscape, it might not make sense to ask for user permission for every website. Second, with the increase of the number of sensors accessible through mobile browsers, this approach might not be effective due to the classic problem of sidestepping the security procedure by users when it is too much of a burden [9]. As stated by one of our participants: “I don’t mind these sensors being risky anyway. I don’t even review the permission list. I have no other choice to be able to use the app”. Moreover,

as we have shown in Section IV, users generally do not understand the implications of these sensors on discovering their PINs for example, even though they know how these sensors work. Hence, such an approach might not be effective in practice.

B. Industrial approach

W3C Device Orientation Event Specification. There is no Security and Privacy section in the latest official W3C Working Draft Document on Device Orientation Event [1]. However, at the time of writing this paper, a new version of the W3C specification is being drafted, which includes a new section on security and privacy issues related to mobile sensors¹⁰, as suggested by us in [22]. The authors working on the revision of the W3C specification point out the problem of fingerprinting mobile devices [7], and touch action recovery [22] through these sensors, and suggest the following mitigations:

- “Do not fire events when the page where they were registered on is not visible or has been backgrounded.”
- “Fire events only on the top-level browsing context or same-origin nested iframes.”
- “Limit the frequency of events (typically 60 Hz seems to be sufficient).”

We believe that these measures may be too restrictive in blocking useful functionalities. For example, imagine a user consciously running a web program in the browser to monitor his daily physical activities such as walking and running. This program needs to continue to have access to the motion and orientation sensor data when the user is working on another tab or minimizes the browser. One might argue that such a program should be available as an app instead, hence the use case is not valid. However, it is expected that the boundary between installed apps and embedded JavaScript programs in the browser will gradually diminish [12].

Mobile browsers. As we showed in [22], browsers and mobile operating systems behave differently on providing access to sensors. Some allow access only on the active webpage and any embedded iframes (although with different origins), some allow access to other tabs, when browser is minimized, or even when the phone is locked. Hence, there is not a consistent approach across all browsers and mobile platforms. Reducing the frequency rate has been applied to all well-known browsers at the moment [22]. For instance, Chrome reduced the sensor readings from 200 Hz to 60 Hz due to security concerns¹¹. However, our attack shows that security risks are still present even at lower frequencies. iOS and Android limit the maximum frequency rate of some sensors such as Gyroscope to 100 Hz and 200 Hz, respectively. It is expected that these frequencies will increase on mobile OSs in the near future and in-browser access is no exception. In fact, current mobile gyroscopes support much higher sampling frequencies, e.g., up to 800 Hz by STMicroelectronics (on Apple products), and up to 8000 Hz by InvenSense (on the Google Nexus range) [23]. With higher frequencies available, attacks such as ours can perform better in the future if adequate security countermeasures are not applied.

Following our report of the issue to Mozilla, starting from version 46 (released in April 2016), Firefox restricts JavaScript access to motion and orientation sensors to only top-level documents and same-origin iframes¹². In the latest Apple Security Updates for iOS 9.3 (released in March 2016), Safari took a similar countermeasure by “suspending the availability of this [motion and orientation] data when the web view is hidden”¹³. However, we believe the implemented countermeasures should only serve as a temporary fix rather than the ultimate solution. In particular, we are concerned that it has the drawback of prohibiting potentially useful web applications in the future. For example, a web page running a fitness program has a legitimate reason to access the motion sensors even when the web page view is hidden. However, this is no longer possible in the new versions of Firefox and Safari. Our concern is confirmed by members in the Google Chromium team¹⁴, who also believe that the issue remains unresolved.

VII. FURTHER DISCUSSION AND LIMITATIONS OF OUR WORK

As mentioned earlier, many of the suggested academic solutions either have not been applied by the industry as a practical solution, or have failed. Given the results in our user studies, designing a practical solution for this problem does not seem to be straightforward. A combination of different approaches might help researchers devise a usable and secure solution. Having control on granting access *before* opening a website and *during* working with it, in combination with a smart notification feature on the browser would probably achieve a balance between security and usability. Users should also have control on reviewing, updating and deleting these data, if stored by the website or shared with a third party *afterwards*. Solutions such as Taintroid [15], a tracking app for monitoring sources of sensitive data on a mobile which has been applied for GPS in [4] could be helpful. After all, it seems that an extensive study is required toward designing a permission framework which is usable and secure at the same time. Such research is a very important usable security and privacy topic to be explored further in the future.

We consider this work a pilot study that explores user risk perception on a comprehensive list of mobile sensors. We envisage the following future work to address these limitations and expand this work:

- *More Participants:* We performed our user studies on a set of users who were recruited from a wide range of backgrounds. Yet the number of the participants is limited. A larger set of participants will improve the confidence in the results. With a large and diverse set of participants, we can also study the effect of demographic factors on perceived risk.
- *Other Risks:* We studied the perceived risk on PINs as a serious and immediate risk to users’ security. The study can be expanded by studying users’ risk perception on other issues such as attackers discovering phone call timing, physical activities, or shopping habits.

¹⁰w3c.github.io/deviceorientation/spec-source-orientation.html

¹¹bugs.chromium.org/p/chromium/issues/detail?id=421691

¹²mozilla.org/en-US/security/advisories/mfsa2016-43/

¹³support.apple.com/en-gb/HT206166

¹⁴bugs.chromium.org/p/chromium/issues/detail?id=523320

- *Other Types of Access:* When interviewing our participants, we presented them with a scenario involving a game app which is installed on their smartphone. This only covers the in-app access to sensors. However, people might express different risk levels for other types of access, e.g., in-browser access. This needs further investigation.
- *Issues with Training Users.* We decided to provide our participants with a short description of each sensor’s functionality (details in the Appendix, part 3). Furthermore, the participants were given the chance to ask as many questions as they wanted to fully understand the functionality of each sensor. This might not be the most effective way to inform users about sensors since some descriptions might seem too technical (and hence not fully understandable) to some users. How to inform users in an effective way is a complex topic of research which can be explored in the future. Besides, we used the same set of participants to generally compare the level of perceived risk before and after seeing sensor descriptions. An alternative approach is to use a different set of participants, i.e., to follow a between-subjects approach instead of a within-subjects one, which would have less bias if carefully designed. However, in order to get meaningful results, the between-subjects approach would require recruiting a larger number of participants.

VIII. CONCLUSION

In this paper, we introduced PINlogger.js, a web-based program which reveals users’ PINs by recording the mobile device’s orientation and motion sensor data through JavaScript code. We also showed that users do not generally perceive a high risk about such sensors being able to steal their PINs. We discussed the complexity of designing a usable and secure solution to prevent the proposed attacks. Access to mobile sensor data via JavaScript is limited to only a few sensors at the moment. This will probably expand in the future, considering for instance the ongoing development of JavaScript-based operating systems such as Firefox OS¹⁵. Hence, designing a general mechanism for secure and usable sensor data management remains a crucial open problem for future research.

IX. ACKNOWLEDGEMENTS

We would like to thank Dr. Kovila Coopamootoo from Newcastle University for her constructive feedback on designing the user studies of this paper. We also would like to thank the voluntary participants who contributed to our data collection and user studies. All our experiments were approved by Newcastle University’s ethical committee. The last three authors are supported by ERC Starting Grant No. 306994.

REFERENCES

[1] W3C Working Draft Document on Device Orientation Event. <http://www.w3.org/TR/orientation-event/>.

¹⁵developer.mozilla.org/en/docs/Mozilla/Firefox_OS/Platform/Architecture

[2] On the best sensor for keystrokes inference attack on android. *Procedia Technology*, 11(0):989 – 995, 2013. 4th International Conference on Electrical Engineering and Informatics, 2013.

[3] A. J. Aviv, B. Sapp, M. Blaze, and J. M. Smith. Practicality of accelerometer side channels on smartphones. In *Proceedings of the 28th Annual Computer Security Applications Conference*, pages 41–50. ACM, 2012.

[4] R. Balebako, J. Jung, W. Lu, L. F. Cranor, and C. Nguyen. “little brothers watching you:” raising awareness of data leaks on smartphones. In *Symposium on Usable Privacy and Security*. ACM Association for Computing Machinery, July 2013.

[5] D. Bichler, G. Stromberg, M. Huemer, and M. Löw. Key generation based on acceleration data of shaking processes. In J. Krumm, G. Abowd, A. Seneviratne, and T. Strang, editors, *UbiComp 2007: Ubiquitous Computing*, volume 4717 of *Lecture Notes in Computer Science*, pages 304–317. Springer Berlin Heidelberg, 2007.

[6] H. Bojinov, Y. Michalevsky, G. Nakibly, and D. Boneh. Mobile device identification via sensor fingerprinting. *CoRR*, abs/1408.1416, 2014.

[7] H. Bojinov, Y. Michalevsky, G. Nakibly, and D. Boneh. Mobile device identification via sensor fingerprinting. *CoRR*, abs/1408.1416, 2014.

[8] J. Bonneau, S. Preibusch, and R. Anderson. A birthday present every eleven wallets? the security of customer-chosen banking pins. In A. Keromytis, editor, *Financial Cryptography and Data Security*, volume 7397 of *Lecture Notes in Computer Science*, pages 25–40. Springer Berlin Heidelberg, 2012.

[9] C. Bravo-Lillo, S. Komanduri, L. F. Cranor, R. W. Reeder, M. Sleeper, J. Downs, and S. Schechter. Your attention please: Designing security-decision uis to make genuine risks harder to ignore. In *Proceedings of the Ninth Symposium on Usable Privacy and Security*, SOUPS ’13, pages 6:1–6:12, New York, NY, USA, 2013. ACM.

[10] L. Cai and H. Chen. Touchlogger: Inferring keystrokes on touch screen from smartphone motion. In *HotSec*, 2011.

[11] L. Cai and H. Chen. On the practicality of motion based keystroke inference attack. In S. Katzenbeisser, E. Weippl, L. Camp, M. Volkamer, M. Reiter, and X. Zhang, editors, *Trust and Trustworthy Computing*, volume 7344 of *Lecture Notes in Computer Science*, pages 273–290. Springer Berlin Heidelberg, 2012.

[12] A. Charland and B. Leroux. Mobile application development: Web vs. native. *Commun. ACM*, 54(5):49–53, May 2011.

[13] A. Das, N. Borisov, and M. Caesar. Exploring ways to mitigate sensor-based smartphone fingerprinting. *CoRR*, abs/1503.01874, 2015.

[14] A. De Luca, A. Hang, E. von Zezschwitz, and H. Hussmann. I feel like i’m taking selfies all day!: Towards understanding biometric authentication on smartphones. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI ’15, pages 1411–1414, New York, NY, USA, 2015. ACM.

[15] W. Enck, P. Gilbert, S. Han, V. Tendulkar, B.-G. Chun, L. P. Cox, J. Jung, P. McDaniel, and A. N. Sheth. Taintdroid: An information-flow tracking system for realtime privacy monitoring on smartphones. *Transactions on Computer Systems.*, June 2014.

[16] Google. Location and sensors apis. Available at: developer.android.com/guide/topics/sensors/index.html.

[17] M. Harbach, E. von Zezschwitz, A. Fichtner, A. D. Luca, and M. Smith. It’s a hard lock life: A field study of smartphone (un)locking behavior and risk perception. In *Symposium On Usable Privacy and Security (SOUPS 2014)*, pages 213–230, Menlo Park, CA, July 2014. USENIX Association.

[18] J. K. Hyungsub Kim, Sangho Lee. Exploring and mitigating privacy threats of html5 geolocation api. In *Annual Computer Security Applications Conference (ACSAC)*. New Orleans, Louisiana, USA, 2014.

[19] R. Kang, L. Dabbish, N. Fruchter, and S. Kiesler. “my data just goes everywhere:” user mental models of the internet and implications for privacy and security. In *Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*, pages 39–52, Ottawa, July 2015. USENIX Association.

[20] H. Khan, U. Hengartner, and D. Vogel. Usability and security perceptions of implicit authentication: Convenient, secure, sometimes annoying. In *Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*, pages 225–239, Ottawa, July 2015. USENIX Association.

- [21] M. Mehrnezhad, E. Toreini, S. F. Shahandashti, and F. Hao. Touchsignatures: Identification of user touch actions based on mobile sensors via javascript (extended abstract). In *Proceedings of the 9th ACM Symposium on Information, Computer and Communications Security (ASIACCS 2015)*. ACM, 2015.
- [22] M. Mehrnezhad, E. Toreini, S. F. Shahandashti, and F. Hao. Touchsignatures: Identification of user touch actions and pins based on mobile sensor data via javascript. *Journal of Information Security and Applications*, 26:23 – 38, 2016.
- [23] Y. Michalevsky, D. Boneh, and G. Nakibly. Gyrophone: Recognizing speech from gyroscope signals. In *Proc. 23rd USENIX Security Symposium*, 2014.
- [24] E. Miluzzo, A. Varshavsky, S. Balakrishnan, and R. R. Choudhury. Tapprints: your finger taps have fingerprints. In *Proceedings of the 10th international conference on Mobile systems, applications, and services*, pages 323–336. ACM, 2012.
- [25] M. F. Mller. A scaled conjugate gradient algorithm for fast supervised learning. *Neural Networks*, 6(4):525 – 533, 1993.
- [26] S. Narain, A. Sanatinia, and G. Noubir. Single-stroke language-agnostic keylogging using stereo-microphones and domain specific machine learning. In *Proceedings of the 2014 ACM Conference on Security and Privacy in Wireless and Mobile Networks*, WiSec '14, pages 201–212, New York, NY, USA, 2014. ACM.
- [27] E. Owusu, J. Han, S. Das, A. Perrig, and J. Zhang. Accessory: password inference using accelerometers on smartphones. In *Proceedings of the Twelfth Workshop on Mobile Computing Systems & Applications*, page 9. ACM, 2012.
- [28] A. Raij, A. Ghosh, S. Kumar, and M. Srivastava. Privacy risks emerging from the adoption of innocuous wearable sensors in the mobile environment. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, pages 11–20, New York, NY, USA, 2011. ACM.
- [29] L. Simon and R. Anderson. Pin skimmer: Inferring pins through the camera and microphone. In *Proceedings of the Third ACM Workshop on Security and Privacy in Smartphones & Mobile Devices*, SPSM '13, pages 67–78, New York, NY, USA, 2013. ACM.
- [30] R. Spreitzer. Pin skimming: Exploiting the ambient-light sensor in mobile devices. In *Proceedings of the 4th ACM Workshop on Security and Privacy in Smartphones & Mobile Devices*, SPSM '14, pages 51–62, New York, NY, USA, 2014. ACM.
- [31] T. Watanabe, M. Akiyama, T. Sakai, and T. Mori. Understanding the inconsistencies between text descriptions and the use of privacy-sensitive resources of mobile apps. In *Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*, pages 241–255, Ottawa, July 2015. USENIX Association.
- [32] K. Y. W. D. H. Y. G. N. P. Xing Jin, Xunchao Hu. Code injection attacks on html5-based mobile apps: Characterization, detection and mitigation. In *Proc. 21th ACM Conference on Computer and Communications Security*, 2014.
- [33] Z. Xu, K. Bai, and S. Zhu. Taplogger: Inferring user inputs on smartphone touchscreens using on-board motion sensors. In *Proceedings of the fifth ACM conference on Security and Privacy in Wireless and Mobile Networks*, pages 113–124. ACM, 2012.

APPENDIX

Interview Script

Hi. Thanks very much for contributing to our study. In this interview, we will ask you to fill in a few questionnaires about mobile sensors such as GPS, camera, light, motion and orientation. You are encouraged to think out loud as you go through, and please feel free to provide any comments during the interview. There is no right or wrong answer, and our purpose is to evaluate the mobile sensors, not you. Everything about this interview is anonymous. Please provide some information about yourself in Table IV.

| | |
|--|--|
| Age | |
| Gender | |
| Profession/ background (optional) | |
| 1st language (optional) | |
| Mobile device | |
| Duration of owning a smartphone/tablet | |

TABLE IV. DEMOGRAPHY

PART ONE

A list of multiple mobile sensors is presented below. To what extent do you know each sensor on a mobile device? Please rate them in the table (Table V was used).

PART TWO

Imagine that you own a smartphone which is equipped with all these sensors. Consider this scenario: you have opened a game app which can have access to all mobile sensors. You leave the game app open in the background, and open your banking app which requires you to enter your PIN.

Do you think any of these sensors can help the game app discover your entered PIN? To what extent are you concerned about each sensor's risk to your PIN? Please rate them in the table (Table VI was used). In this section, please only rely on the knowledge you already have about the sensors, and if you do not know some of them, describe your feeling of security about them.

PART THREE

Let us explain each sensor here:

- GPS: identifies the real-world geographic location.
- Camera, Microphone: capture pictures/videos and voice, respectively.
- Fingerprint, TouchID: scans the fingerprint.
- Touch Screen: enables the user to interact directly with the display by physically touching it.
- WiFi: is a wireless technology that allows the device to connect to a network.
- Bluetooth: is a wireless technology for exchanging data over short distances.
- NFC (Near Field Communication): is a wireless technology for exchanging data over shorter distances (less than 10 cm) for purposes such as contactless payment.
- Proximity: measures the distance of objects from the touch screen.
- Ambient Light: measures the light level in the environment of the device.
- Ambient Pressure (Barometer), Ambient Humidity, and Ambient Temperature: measure the air pressure, humidity, and temperature in the environment of the device, respectively.
- Device Temperature: measures the temperature of the device.
- Gravity: measures the force of gravity.

- Magnetic Field: reports the ambient magnetic field intensity around the device.
- Hall sensor: produces voltage based on the magnetic field.
- Accelerometer: measures the acceleration of the device movement or vibration.
- Rotation: reports how much and in what direction the device is rotated.
- Gyroscope: estimates the rotation rate of the device.
- Motion: measures the acceleration and the rotation of the device.
- Orientation: reports the physical angle that the device is held in.
- Sensor Hub: is an activity recognition sensor and its purpose is to monitor the device's movement.

Please feel free to ask us about any of these sensors for more information.

Now that you have more knowledge about the sensors, let us describe the same scenario here again. Imagine that you own a smartphone which is equipped with all these sensors. You have opened a game app which can have access to all mobile sensors. You leave the game app open in the background, and open your banking app which requires you to enter your PIN.

Do you think any of these sensors can help the game app to discover your entered PIN? To what extent are you concerned about each sensor's risk to your PIN? Please rate them in the table (Table VI was used). In this part, please make sure that you know the functionality of all the sensors. If you are unsure, please have another look at the descriptions, or ask us about them.

Thanks very much for taking part in this study. Please leave any extra comment here.

An Amazon voucher and a business card are in this envelope. Please contact us if you have any questions about this interview, or are interested in the results of this study.

| Sensor | I've never heard of this | I've heard of this but I don't know what this is | I know what this is but I don't know how this works | I know generally how this works | I know very well how this works |
|---------------------|--------------------------|--|---|---------------------------------|---------------------------------|
| Bluetooth | | | | | |
| Gyroscope | | | | | |
| GPS | | | | | |
| Sensor Hub | | | | | |
| Ambient Temperature | | | | | |
| Accelerometer | | | | | |
| Magnetic Field | | | | | |
| Motion | | | | | |
| Fingerprint | | | | | |
| Orientation | | | | | |
| Proximity | | | | | |
| Ambient Pressure | | | | | |
| Hall Sensor | | | | | |
| Rotation | | | | | |
| Touch Screen | | | | | |
| Camera | | | | | |
| TouchID | | | | | |
| Barometer | | | | | |
| Gravity | | | | | |
| Microphone | | | | | |
| Ambient Humidity | | | | | |
| WiFi | | | | | |
| Ambient Light | | | | | |
| NFC | | | | | |
| Device Temperature | | | | | |

TABLE V. THIS FORM WAS USED FOR PART ONE

| Sensor | Risk to PIN | | | | |
|---------------------|---------------|--------------------|----------------------|-----------|---------------------|
| | Not Concerned | A little Concerned | Moderately Concerned | Concerned | Extremely Concerned |
| Bluetooth | | | | | |
| Gyroscope | | | | | |
| GPS | | | | | |
| Sensor Hub | | | | | |
| Ambient Temperature | | | | | |
| Accelerometer | | | | | |
| Magnetic Field | | | | | |
| Motion | | | | | |
| Fingerprint | | | | | |
| Orientation | | | | | |
| Proximity | | | | | |
| Ambient Pressure | | | | | |
| Hall Sensor | | | | | |
| Rotation | | | | | |
| Touch Screen | | | | | |
| Camera | | | | | |
| TouchID | | | | | |
| Barometer | | | | | |
| Gravity | | | | | |
| Microphone | | | | | |
| Ambient Humidity | | | | | |
| WiFi | | | | | |
| Ambient Light | | | | | |
| NFC | | | | | |
| Device Temperature | | | | | |

TABLE VI. THIS FORM WAS USED FOR PARTS TWO AND THREE