



Unexpected Inferences from Sensor Data: A Hidden Privacy Threat in the Internet of Things

Jacob Kröger

► To cite this version:

Jacob Kröger. Unexpected Inferences from Sensor Data: A Hidden Privacy Threat in the Internet of Things. 1st IFIP International Internet of Things Conference (IFIPIoT), Sep 2018, Poznan, Poland. pp.147-159, 10.1007/978-3-030-15651-0_13 . hal-03217368

HAL Id: hal-03217368

<https://inria.hal.science/hal-03217368>

Submitted on 4 May 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Unexpected Inferences from Sensor Data: A Hidden Privacy Threat in the Internet of Things

Jacob Kröger

Technische Universität Berlin
Straße des 17. Juni 135, 10623, Berlin, Germany
kroeger@tu-berlin.de

Abstract. A growing number of sensors, embedded in wearables, smart electric meters and other connected devices, is surrounding us and reaching ever deeper into our private lives. While some sensors are commonly regarded as privacy-sensitive and always require user permission to be activated, others are less protected and less worried about. However, experimental research findings indicate that many seemingly innocuous sensors can be exploited to infer highly sensitive information about people in their vicinity. This paper reviews existing evidence from the literature and discusses potential implications for consumer privacy. Specifically, the analysis reveals that certain insufficiently protected sensors in smart devices allow inferences about users' locations, activities and real identities, as well as about their keyboard and touchscreen inputs. The presented findings call into question the adequacy of current sensor access policies. It is argued that most data captured by smart consumer devices should be classified as highly sensitive by default. An introductory overview of sensors commonly found in these devices is also provided, along with a proposed classification scheme.

Keywords: Privacy, Sensors, Internet of Things, Inference Attacks, Information Leaks

1 Introduction

At the latest since the advent of smartphones, a variety of embedded sensors is constantly surrounding us – whether we are at work, in transit, or even within our own four walls. The emerging Internet of Things (IoT) is predicted to further increase the number of sensors in our everyday environment by many orders of magnitude [17].

New services and business models are enabled through the increasing pervasiveness and interconnection of sensors which promise to bring transformational improvements in many areas including health, safety, security, convenience, productivity and sustainability. At the same time, with sensing technologies reaching ever deeper into people’s lives, there is growing concern about potential privacy violations. These concerns need to be addressed before IoT technologies are widely deployed – not only to protect the fundamental right to informational self-determination, but also to foster trust and acceptance among users.

While cameras, microphones and navigation systems like GPS are commonly perceived as privacy-sensitive [26, 32] and require explicit user permission in current mobile operating systems [6], many inconspicuous sensors such as accelerometers, gyroscopes and barometers are less well-understood in terms of their privacy implications, and also less protected [57]. Countless embedded sensors in consumer electronics, such as smartphones, smartwatches, and smart home appliances, can be freely accessed by various possibly untrusted parties ranging from device manufacturers [48] and service providers [59] to third-party apps installed on mobile devices [6] and even website operators [8].

Numerous studies have shown that highly personal information can be inferred from seemingly innocuous sensor data. Accordingly, an international group of data protection and privacy commissioners argued in the *Mauritius Declaration on the Internet of Things* that all IoT sensor data should be treated as personal data [33]. This paper substantiates this view with empirical evidence from the existing literature and illustrates potential implications for consumer privacy.

2 Sensors in IoT Devices

2.1 Classification

A wide variety of sensors can be found in consumer IoT devices. Depending on the type of device and its field of application (e.g. home automation, fitness monitoring, gaming), sensor measurements fulfil many different purposes. In the context of IoT consumer privacy, sensors can be meaningfully classified according to the following properties:

Measured Variable. Sensors can be classified based on the specific physical or chemical property they measure (e.g. air pressure, temperature, acceleration), or – in the case of fusion sensors which combine the results of multiple underlying sensors – the computed result they provide (e.g. absolute orientation of a device).

In the privacy context, this seems more accurate and meaningful than the widespread classification of sensors based on their intended purpose. In Android, for instance, accelerometers are exclusively classified as motion sensors because their declared purpose is to monitor a device's motion [1]. More relevant for assessing their potential privacy impact, however, is the fact that accelerometers can also be used to infer much more sensitive information, including a user's location [24, 27], activities [55], and login credentials [49]. Therefore, a classification of sensors based solely on their most common purpose can be misleading.

Internal or External Focus. It can be differentiated between sensors which measure internal properties of their encapsulating device (e.g. internal temperature) and sensors which measure environmental properties (e.g. ambient illumination).

Possible Deployments. Depending on their size, functionality and intended purpose, sensors can vary in mobility. While some sensors are usually stationary (e.g. carbon dioxide sensors) or even wall-mounted (e.g. passive infrared motion detectors), others can be carried around or even be embedded into mobile devices (e.g. microphones, gyroscopes, accelerometers).

Reporting Mode. Sensors can be classified based on how frequently they generate events. Android, for instance, differentiates between continuous reporting (events are generated at a constant rate), on-change reporting (only if the measured values have changed), one-shot reporting (upon detection of a specific event) and special reporting (according to specific sensor requirements) [2].

Level of Protection. Depending on the assumed sensitivity of their data and other factors, sensors in IoT devices enjoy different levels of legal and technical protection. Specific protection mechanisms and access policies can vary significantly between devices of different manufacturers [6]. Further insights into the sensitivity and protection of different types of sensor data are provided in section 2.2.

Accuracy. The measurement accuracy of sensors can greatly differ and often depends on their technological complexity and production quality. Sometimes, the granularity of sensor data is intentionally reduced to enhance privacy [16].

Visibility. Sensors can also be classified based on their degree of visibility. While some sensors can be completely hidden within a device (e.g. accelerometer) or even sense through walls (e.g. radio tomography), other sensors require visual contact to function (e.g. infrared). Also, the size and appearance of a sensor should be considered when assessing its degree of visibility.

2.2 Sensitivity and Protection of Sensor Data

There are considerable differences in the perceived privacy intrusiveness of different sensor types. Cameras, for instance, are commonly regarded as highly privacy-sensitive [34] and therefore often rejected in private rooms [44]. Similarly, microphones [13], biometrical sensors [44], and GPS data [32] are widely believed to intrude user privacy. By contrast, sensors that do not directly record or locate the user are often considered harmless. Examples include inertial sensors such as accelerometers [32] and sensors capturing environmental data such as room temperature [44].

This sentiment is reflected in the permission policies of current mobile operating systems. All mainstream platforms, including iOS, Android and Windows Phone, allow third-party apps to access various built-in sensors without requiring security permission [6]. An overview of common smartphone sensors and their permission requirements on Android is provided in Table 1. While Google’s mobile operating system leaves numerous sensors unprotected, access to GPS, camera and microphone is restricted [38].

It can be expected that the policies of technology and market leaders like Apple, Google and Microsoft will substantially inspire and shape data protection mechanisms in the future IoT. Apart from mobile devices, other IoT devices such as smart meters [16, 23, 59], motion detectors [25, 59] and smoke alarm systems [48] have also been shown to share potentially sensitive data with service providers and device manufacturers without privacy-related restrictions. It is therefore important to investigate whether the widely assumed innocuousness of certain sensors matches reality.

Table 1. Overview of common smartphone sensors

Sensor type	Output	Restricted in Android [6, 38]
Accelerometer	Acceleration of device	No
Ambient Temperature	Room / air temperature	No
Camera	Recorded image or video	Yes
GPS	Geographical location	Yes
Gravity	Force of gravity	No
Gyroscope	Device’s rate of rotation	No
Light	Level of ambient illumination	No
Magnetometer	Ambient geomagnetic field	No
Microphone	Recorded audio	Yes
Orientation	Device’s degrees of rotation	No
Pressure	Ambient air pressure	No
Proximity	Proximity of an object to the device’s screen	No
Relative Humidity	Relative ambient air humidity	No
Temperature	Temperature of the device	No

3 Localization

3.1 Indoor Localization and Occupancy Monitoring

Various unobtrusive sensors can be used for the purpose of monitoring room occupancy and locating people indoors. A technology commonly found in home security systems are passive infrared (PIR) sensors which detect moving objects based on changes in ambient temperature. Besides from binary room occupancy, PIR sensors have been used to infer the number of occupants and even to identify people [59].

Since humans naturally exhale carbon dioxide (CO_2), the air composition in a room or building can also be an indicator of occupancy. CO_2 sensors have already been exploited to estimate the number of room occupants with high accuracy [3]. Similarly, human presence can be detected through measuring other environmental parameters, such as ambient humidity, illumination, and sound rate [4].

Even the on/off state of household appliances and their level of power consumption over time allow inferences about binary [16] and ranged [59] room occupancy. The energy usage patterns required for such inferences can be derived from residential smart meter readings [18] which are commonly shared with electric utilities for the official purpose of demand-response management and billing [59].

Radio Tomographic Imaging is another technique that has been exploited for human presence detection. This approach is based on the fact that the human body absorbs and thereby weakens a radio signal between sender and receiver, which can be measured in terms of the received signal strength (RSS). Within wireless sensor networks, which are considered an enabling technology for IoT [52], radio tomographic imaging can be used for indoor localization [10] and even multi-target tracking [45] without requiring any additional hardware.

Radar sensors have also been proposed for the purpose of indoor localization. Aside from just detecting human presence and motion, existing approaches can sense through walls [21] and investigate the physical characteristics of targets [31]. Other technologies that are being explored for unobtrusive indoor positioning and room-level tracking include ultrasonic sensors [26], door-operated switches [30], and pressure-sensing floors [51]. In contrast to some of the data sources mentioned earlier, sensors like radar and ultrasound are unlikely to be underestimated in terms of their sensitivity. However, these technologies may be completely invisible to the tracked person which could be utilized for covert intrusions of privacy.

Indoor location information can be highly privacy-sensitive. Among many other possible inferences, indoor location traces have been used to obtain personal properties of targets, such as age, work role, coffee consumption and smoking habits [41].

3.2 Geographical Localization

Studies have shown that sensors in mobile devices can be exploited for geographical user localization, even when GPS is disabled. For instance, accelerometers in smartphones can reveal the device's location while the user is driving a car [24] or using a metropolitan train system [27]. These approaches are either based on supervised

learning with labeled training data or on mapping the phone’s motion trajectory to the shape of existing routes on a map through computing a similarity score.

Smartphone owners can also be located while they are walking or travelling on a train or plane by exploiting a combination of easily accessible mobile sensor data and publicly available auxiliary data [43]. For this approach, neither training data nor prior knowledge about the user is required. Researchers have even shown that a smartphone’s location can be learned from tracking its battery level over time [42]. This method builds on the fact that the level of power consumption is affected by the distance and any obstacles between a phone and its current cellular base station.

Highly sensitive information can be inferred from location traces. Through spatial clustering and the identification of temporal mobility patterns, it is possible to obtain the true identity as well as the home address and points of interest of a person from his or her location history [19]. Spatio-temporal information can also be used to infer a person’s mode of transportation and to predict where a driver will drive [35]. It is important to understand that certain visited locations alone (e.g. betting office, mosque, vegetarian restaurant) can reveal sensitive personal properties such as habits, preferences or religious affiliation. In combination with various unprotected smartphone sensors, location data has even been used to draw inferences about a user’s emotional state [37].

4 Activity Recognition

Data from various inconspicuous sensors in IoT devices can be used to infer information about basic physical activities (e.g. walking, running, climbing) and activities of daily living (e.g. grooming, cooking, washing dishes) of individuals in their vicinities.

Recent studies in the field of sensor-based human activity recognition have mostly focused on accelerometer data from wearable devices such as smartphones, smartwatches, or wrist-worn fitness trackers [11]. In [55], 3-axis accelerometry collected with an off-the-shelf smartwatch was successfully used to recognize eating moments of users in free-living conditions. Data from wrist-worn inertial sensors has also been used to recognize various other actions including drinking and cigarette smoking [50, 54]. Whole-body motions, such as descending and ascending stairs, walking, jumping, and jogging, can be recognized with high accuracy from smartphone inertial sensor data [11]. Motion sensors in smartphones have even been used for the assessment of driving behavior [29].

Another unobtrusive but revealing data source for activity recognition in an IoT environment are smart electric meters. Various behaviors of residents including bathroom activities, cooking, housework, sleep cycles, and meal times can be inferred from seemingly non-sensitive smart meter readings [16]. It has been shown that even the current TV channel and specific audiovisual content displayed on a television can be identified based on the corresponding household’s electricity usage profile [23].

Data from infrared motion detectors, which is normally used for monitoring room occupancy and often shared with security companies [59], has also been exploited for

activity recognition. In [25], passive infrared sensors were used to detect the operation mode of kitchen appliances, such as water cooker and toaster, as well as human actions, such as opening a refrigerator and taking a shower. While many current techniques for automated activity recognition require pre-processing of collected data, approaches have already been presented for extracting activity-related information from infrared sensor data in real-time [34].

In a smart home environment, data from multiple sources can be combined for improved recognition accuracy. Nef et al. [47] used a set of sensors measuring several parameters including luminescence, temperature and ambient motion to recognize eight different activities of daily living such as grooming, cooking, eating, and watching TV.

5 User Input Inference

Among the most sensitive information in the context of consumer electronics is the input that users provide through touchscreens and keyboards, which includes personal notes, text messages, transaction details, and login credentials.

Multiple so-called *keystroke inference attacks* have been developed for inferring such user inputs from seemingly innocuous sensor data. Most of these attacks are based on the observation that the keys typed into a device correlate with micro-motions of the user's hand, and – in the case of mobile devices – with micro-motions of the device itself. Owusu et al. [49] show that data from accelerometers in smartphones can be used to obtain entire sequences of text entered through the touchscreen, including passwords. The same kind of data can also be exploited to infer graphical password patterns [5]. As demonstrated by Spreitzer [53], even ambient-light sensors in mobile devices may leak sensitive user inputs such as the personal identification number (PIN).

Not only can embedded sensors reveal user inputs to their encapsulating device, they can also allow inferences about what users type into other devices. Motion sensors in modern smartwatches, for instance, have been exploited to recover text typed on laptop and desktop keyboards [40, 56].

The possibility to infer user inputs from sensor data could enable malicious parties to gain insight into sensitive communications and transactions. Beltramelli [7] even suggests that a user's entire technological ecosystem could be compromised when passwords are leaked through embedded sensors in consumer IoT devices. Although current keystroke inference attacks are mostly based on accelerometer and gyroscope data, it has been proposed that other potential side channels such as magnetic field sensors should also be investigated [56].

6 Identification

6.1 De-anonymization of Sensor Data

Sensor data collected through consumer IoT devices is particularly sensitive when it can be linked to the real identity of the user. In scenarios where data is accessed by

untrusted service providers or device manufacturers, identifying information such as name, address and bank account is often already available to them from the respective service or purchase agreement.

Even if this is not the case, the sensor data itself and inferred location information can often reveal a user's identity to the data controller. As detailed above, work and home addresses of individuals can be derived from their geographical location traces [19]. Such information – even at coarse resolution – can be used in conjunction with employment directories, white pages, tax records, and other public or private datasets to uniquely identify a tracked person [22]. Social networks have also been proven to be excellent sources of auxiliary data for the de-anonymization of smartphone sensor data [36]. Some researchers even argue that no method exists to reliably prevent the de-anonymization of location data [46].

The linkage of data sources has been recognized as one of the major privacy threats in IoT [60]. The term *quasi-identifier*, coined by Dalenius [14], is generally used to describe pieces of information that are not personally identifiable when considered in isolation but can quickly become unique identifiers when combined with other data sources. When it comes to IoT sensor data, an adversary's auxiliary information does not even need to be drawn from external sources (e.g. information brokers, public datasets) but can also be collected from within a user's own ecosystem of connected devices. As shown in [36], mobile sensor data can become personally identifiable through linkage with less protected data streams of the same user. Seemingly innocuous sensors could thus be used as gateways to de-anonymize sensitive streams of user activity.

6.2 Tracking and Profiling of Anonymous Users

Even in cases where an adversary is not capable of obtaining a target's real identity from collected IoT sensor data, substantial privacy issues may arise. For some of the most common intentions behind data misuse, such as price discrimination and targeted advertising, it can be fully sufficient to distinguish between anonymous individuals without ever learning their real names.

In the context of consumer IoT devices, unidentified users can be told apart based on their physical characteristics and movement patterns, which are often reflected in sensor data from mobile devices. According to Jain et al. [28], body movements can serve as biological fingerprints if they are universal, repeatable, collectible, and distinctive. It has been shown that head movements, as captured through the sensors of head-worn IoT devices, fulfill these requirements [39]. Similarly, users of sensing devices can be uniquely distinguished based on their typing motion behavior [20] and other natural gestures such as lowering/raising an arm [58]. The required motion data can be drawn from sensors in smartphones [9] and wrist-worn devices [58] without requiring any permission or conscious participation from the user. Where accessible through IoT sensors, physical characteristics such as body weight may also be used to distinguish observation targets from one another [26].

Following another approach called *device fingerprinting*, users can be told apart based on imperfections and specific features of their personal devices. Even calibration

errors in sensors commonly regarded as non-sensitive, such as gyroscopes and accelerometers, have been exploited to uniquely identify IoT devices [8, 15].

The ability to recognize and track individuals through sensor data is a potential basis for behavioral profiling [60] and can thus result in people being treated differently, even without names attached to the data .

7 Discussion and Implications

As shown in the previous sections, sensitive information regarding a user’s location, activities and device inputs can be obtained from seemingly benign IoT sensor data. Through cross-linking different sources of sensor data with auxiliary information, an adversary can potentially identify the victim and draw further undesired inferences about his or her habits and preferences. Thus, even if an IoT device is not designed or expected to capture sensitive information, its data streams can indirectly enable serious invasions of user privacy.

It can be assumed that this threat will continue to grow with further improvements of sensor technologies in terms of size, cost and accuracy, further advances in machine learning methods, and – most importantly – the predicted rapid proliferation of consumer IoT devices [17]. Therefore, the privacy-intrusion potential of seemingly innocuous sensors needs to be addressed and dealt with urgently in order to effectively protect consumers’ fundamental right to privacy.

The level of legal and technical protection of sensor data should always be chosen in consideration of all reasonably conceivable inferences that can be derived from the data, and not based on the sensor’s official purpose. No distinct line can be drawn, for instance, between GPS data and currently less protected sensors such as accelerometers and gyroscopes when it comes to their potential of revealing sensitive information [24, 27]. Further research into the privacy impact of specific sensors is needed, taking into account state-of-the-art data mining and machine learning techniques. As it is extremely difficult, however, to meaningfully determine the limits of continuously advancing inference attacks, most types of IoT sensor data should be classified as highly sensitive by default.

Consumers need to be warned about the invisible privacy threats of IoT. Although surveys have shown that many individuals do not want IoT devices to record all of their inhome behaviors [12] and worry about inferences from otherwise anonymous data [44], people are not particularly concerned about unobtrusive sensors such as the ones discussed in this paper [32, 44]. It seems apparent that the possibilities and implications of data linkage and pattern recognition are not well-understood by the average consumer. Thus, truly informed consent to data processing, as required by the EU’s General Data Protection Regulation and other data protection laws, can often not be obtained. Improved education of consumers and their voiced objection against the insufficient protection of sensitive sensor data could also exert pressure on policy makers and device manufacturers to incorporate effective privacy preserving mechanisms in IoT systems and the corresponding regulation.

In order to enable consumers to remain in control of their personal data and understand the complex privacy implications of IoT sensor technology, information on data collection and processing must be presented in an intuitive and easily understandable manner. For this purpose, new tools and approaches need to be developed.

8 Conclusion

The emerging Internet of Things promises to improve our quality of life in many ways. However, the large variety of sensors embedded into connected devices also poses a rising threat to consumer privacy. The overview provided in this paper illustrates that even sensors that are generally thought to be harmless can be used to infer highly sensitive personal information, such as a user's location, activities and real identity, as well as keyboard and touchscreen inputs. Since the access to seemingly innocuous sensors is often unrestricted, various parties can regularly capture and use their potentially revealing data, including device manufacturers, service providers, and third-party apps installed on mobile devices.

Many sensors discussed in this paper are usually hidden inside their encapsulating device and thus completely invisible to the user. Where no access permission is required, these sensors might enable fully covert surveillance of device owners and other people in their vicinity. Considering the extensive protection of sensors which are commonly believed to be privacy-sensitive, such as cameras, microphones and GPS, it is important to prevent malicious parties from using less-protected sensors as substitutional data sources.

Judging from the research findings presented in this paper, many disregarded types of sensor data in the IoT should be classified as sensitive by default and protected accordingly. Furthermore, improved consumer education and more intuitive ways of presenting privacy-related information are necessary to achieve true informational self-determination in an increasingly complex and diverse environment of sensing devices. This paper, however, provides only an initial exploration of the topic. Further empirical and conceptual research into the privacy implications of sensors in IoT devices is strongly encouraged.

References

1. Android Developers: Motion Sensors, https://developer.android.com/guide/topics/sensors/sensors_motion.html.
2. Android Open Source Project: Reporting modes, <https://source.android.com/devices/sensors/report-modes>.
3. Ang, I.B.A. et al.: CD-HOC: Indoor Human Occupancy Counting using Carbon Dioxide Sensor Data. ArXiv170605286 CsHC. (2017).
4. Ang, I.B.A. et al.: Human occupancy recognition with multivariate ambient sensors. In: Pervasive Computing and Communication Workshops (PerCom Workshops). pp. 1–6 (2016).

5. Aviv, A.J. et al.: Practicality of accelerometer side channels on smartphones. In: Proceedings of the 28th Annual Computer Security Applications Conference. pp. 41–50 ACM (2012).
6. Bai, X. et al.: Sensor Guardian: prevent privacy inference on Android sensors. *EURASIP J. Inf. Secur.* 2017, 1, (2017).
7. Beltramelli, T., Risi, S.: Deep-Spying: Spying using Smartwatch and Deep Learning. *ArXiv151205616 Cs.* (2015).
8. Bojinov, H. et al.: Mobile device identification via sensor fingerprinting. *arXiv:1408.1416.* (2014).
9. Buriro, A. et al.: Please hold on: Unobtrusive user authentication using smartphone’s built-in sensors. In: Identity, Security and Behavior Analysis (ISBA). pp. 1–8 IEEE (2017).
10. Cao, X. et al.: A Lightweight Robust Indoor Radio Tomographic Imaging Method in Wireless Sensor Networks. *Prog. Electromagn. Res.* 60, 19–31 (2017).
11. Chen, Y., Shen, C.: Performance Analysis of Smartphone-Sensor Behavior for Human Activity Recognition. *IEEE Access.* 5, 3095–3110 (2017).
12. Choe, E.K. et al.: Living in a glass house: a survey of private moments in the home. In: Proceedings of the 13th International Conference on Ubiquitous Computing. p. 41 ACM Press (2011).
13. Cook, D.J., Krishnan, N.: Mining the Home Environment. *J. Intell. Inf. Syst.* 43, 3, 503–519 (2014).
14. Dalenius, T.: Finding a Needle In a Haystack or Identifying Anonymous Census Records. *J. Off. Stat.* (1986).
15. Das, A. et al.: Tracking Mobile Web Users Through Motion Sensors: Attacks and Defenses. In: Network and Distributed System Security Symposium (NDSS). Internet Society (2016).
16. Eibl, G., Engel, D.: Influence of Data Granularity on Smart Meter Privacy. *IEEE Trans. Smart Grid.* 6, 2, 930–939 (2015).
17. Evans, D.: The Internet of Things - How the Next Evolution of the Internet Is Changing Everything, https://www.cisco.com/c/dam/en_us/about/ac79/docs/in-nov/IoT_IBSG_0411FINAL.pdf.
18. Fan, J. et al.: Privacy Disclosure Through Smart Meters: Reactive Power Based Attack and Defense. In: Dependable Systems and Networks (DSN). pp. 13–24 IEEE (2017).
19. Freudiger, J. et al.: Evaluating the privacy risk of location-based services. In: International conference on financial cryptography and data security. pp. 31–46 Springer (2011).
20. Gascon, H. et al.: Continuous Authentication on Mobile Devices by Analysis of Typing Motion Behavior. In: Proc. of GI Conference “Sicherheit.” pp. 1–12 (2014).
21. Gennarelli, G. et al.: Real-Time Through-Wall Situation Awareness Using a Microwave Doppler Radar Sensor. *Remote Sens.* 8, 8, 621 (2016).
22. Golle, P., Partridge, K.: On the anonymity of home/work location pairs. In: International Conference on Pervasive Computing. pp. 390–397 Springer (2009).
23. Greveler, U. et al.: Multimedia content identification through smart meter power usage profiles. In: Proceedings of the International Conference on Information and Knowledge Engineering (IKE). p. 1 WorldComp (2012).
24. Han, J. et al.: Accomplice: Location inference using accelerometers on smartphones. In: Communication Systems and Networks (COMSNETS). pp. 1–9 IEEE (2012).

25. Hevesi, P. et al.: Monitoring Household Activities and User Location with a Cheap, Unobtrusive Thermal Sensor Array. In: Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing. pp. 141–145 ACM, New York, NY, USA (2014).
26. Hnat, T.W. et al.: Doorjamb: unobtrusive room-level tracking of people in homes using doorway sensors. In: Proceedings of the 10th ACM Conference on Embedded Network Sensor Systems. pp. 309–322 ACM (2012).
27. Hua, J. et al.: We Can Track You If You Take the Metro: Tracking Metro Riders Using Accelerometers on Smartphones. *IEEE Trans. Inf. Forensics Secur.* (2017).
28. Jain, A.K. et al.: An introduction to biometric recognition. *IEEE Trans. Circuits Syst. Video Technol.* 14, 1, 4–20 (2004).
29. Júnior, J.F. et al.: Driver behavior profiling: An investigation with different smartphone sensors and machine learning. *PLOS ONE*. 12, 4, e0174959 (2017).
30. Kim, S.H. et al.: Improved occupancy detection accuracy using PIR and door sensors for a smart thermostat. In: Proceedings of the 15th IBPSA Conference. pp. 2753–2758, San Francisco, CA, USA (2017).
31. Kim, Y. et al.: Human Detection Using Doppler Radar Based on Physical Characteristics of Targets. *IEEE Geosci. Remote Sens. Lett.* 12, 2, 289–293 (2015).
32. Klasnja, P. et al.: Exploring privacy concerns about personal sensing. In: International Conference on Pervasive Computing. pp. 176–183 Springer (2009).
33. Kohnstamm, J., Madhub, D.: Mauritius Declaration on the Internet of Things, https://edps.europa.eu/sites/edp/files/publication/14-10-14_mauritius_declaration_en.pdf.
34. Krishnan, N.C., Cook, D.J.: Activity Recognition on Streaming Sensor Data. *Pervasive Mob. Comput.* 10, Pt B, 138–154 (2014).
35. Krumm, J.: Inference attacks on location tracks. In: International Conference on Pervasive Computing. pp. 127–143 Springer (2007).
36. Lane, N.D. et al.: On the feasibility of user de-anonymization from shared mobile sensor data. In: Proceedings of the Third International Workshop on Sensing Applications on Mobile Phones. p. 3 ACM (2012).
37. Lee, H. et al.: Towards unobtrusive emotion recognition for affective social communication. In: Consumer Communications and Networking Conference (CCNC). pp. 260–264 IEEE (2012).
38. Lee, W.-H., Lee, R.: Multi-sensor authentication to improve smartphone security. In: 2015 International Conference on Information Systems Security and Privacy (ICISSP). (2015).
39. Li, S. et al.: Whose move is it anyway? Authenticating smart wearable devices using unique head movement patterns. In: Pervasive Computing and Communications (PerCom). pp. 1–9 IEEE (2016).
40. Maiti, A. et al.: Smartwatch-Based Keystroke Inference Attacks and Context-Aware Protection Mechanisms. In: Proceedings of the 11th ACM on Asia Conference on Computer and Communications Security. pp. 795–806 ACM, New York, NY, USA (2016).
41. Matsuo, Y. et al.: Inferring Long-term User Properties Based on Users’ Location History. In: *IJCAI*. pp. 2159–2165 (2007).
42. Michalevsky, Y. et al.: PowerSpy: Location Tracking using Mobile Device Power Analysis. In: *USENIX Security Symposium*. (2015).

43. Mosenia, A. et al.: PinMe: Tracking a Smartphone User around the World. *IEEE Trans. Multi-Scale Comput. Syst.* 1–1 (2017).
44. Naeini, P.E. et al.: Privacy Expectations and Preferences in an IoT World. In: *Thirteenth Symposium on Usable Privacy and Security (SOUPS)*. (2017).
45. Nannuru, S. et al.: Radio-frequency tomography for passive indoor multitarget tracking. *IEEE Trans. Mob. Comput.* 12, 12, 2322–2333 (2013).
46. Narayanan, A., Felten, E.W.: No silver bullet: De-identification still doesn't work, <http://randomwalker.info/publications/no-silver-bullet-de-identification.pdf>.
47. Nef, T. et al.: Evaluation of Three State-of-the-Art Classifiers for Recognition of Activities of Daily Living from Smart Home Ambient Data. *Sensors*. 15, 5, 11725–11740 (2015).
48. Notra, S. et al.: An experimental study of security and privacy risks with emerging household appliances. In: *Communications and Network Security (CNS)*. pp. 79–84 IEEE (2014).
49. Owusu, E. et al.: ACCessory: password inference using accelerometers on smartphones. In: *Proceedings of the Twelfth Workshop on Mobile Computing Systems & Applications*. p. 9 ACM (2012).
50. Parate, A.: *Designing Efficient and Accurate Behavior-Aware Mobile Systems*. University of Massachusetts Amherst (2014).
51. Shen, Y.-L., Shin, C.: Distributed Sensing Floor for an Intelligent Environment. *Sens. J. IEEE*. 9, 1673–1678 (2010).
52. Shukri, S., Kamarudin, L.M.: Device free localization technology for human detection and counting with RF sensor networks: A review. *J. Netw. Comput. Appl.* 97, 157–174 (2017).
53. Spreitzer, R.: 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*. ACM, New York, NY, USA (2014).
54. Tang, Q.: *Automated Detection of Puffing and Smoking with Wrist Accelerometers*. Northeastern University Boston (2014).
55. Thomaz, E. et al.: A practical approach for recognizing eating moments with wrist-mounted inertial sensing. In: *Proceedings of the ACM International Conference on Ubiquitous Computing*. pp. 1029–1040 ACM Press (2015).
56. Wang, H. et al.: MoLe: Motion Leaks through Smartwatch Sensors. In: *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*. pp. 155–166 ACM Press (2015).
57. Xu, Z., Zhu, S.: SemaDroid: A Privacy-Aware Sensor Management Framework for Smartphones. In: *Proceedings of the 5th ACM Conference on Data and Application Security and Privacy*. pp. 61–72 ACM Press (2015).
58. Yang, J. et al.: MotionAuth: Motion-based authentication for wrist worn smart devices. In: *Pervasive Computing and Communication Workshops (PerCom Workshops)*. pp. 550–555 IEEE (2015).
59. Yang, L. et al.: Inferring occupancy from opportunistically available sensor data. In: *Pervasive Computing and Communications (PerCom)*. pp. 60–68 IEEE (2014).
60. Ziegeldorf, J.H. et al.: Privacy in the Internet of Things: threats and challenges. *Secur. Commun. Netw.* 7, 12, 2728–2742 (2014).