

A report on personally identifiable sensor data from smartphone devices

Marios Fanourakis
CUI, Quality of Life Lab
University of Geneva, Switzerland

Abstract

An average smartphone is equipped with an abundance of sensors to provide a variety of vital functionalities and conveniences. The data from these sensors can be collected in order to find trends or discover interesting correlations in the data but can also be used by nefarious entities for the purpose of revealing the identity of the persons who generated this data. In this paper, we seek to identify what types of sensor data can be collected on a smartphone and which of those types can pose a threat to user privacy by looking into the hardware capabilities of modern smartphone devices and how smartphone data is used in the literature. We then summarize some implications that this information could have on the GDPR.

1 Introduction

Data can be obtained through many types of sensors or surveys. Already, the average mobile device includes several sensors as a standard feature: accelerometer (to know when the screen rotates), compass (for positioning), GPS (for positioning), light (to adjust the display brightness), audio (microphone), image (camera), and others. Other devices may include different types of sensors like temperature, air quality, heart rate, etc. Moreover, a mobile phone with such sensors roams with its owner, and can be used to collect context information on their behalf.

An average smartphone is equipped with an abundance of sensors to provide a variety of vital functionalities and conveniences. For example, the basic telephony antenna which enables the smartphone to connect to the cellular network, or the ambient light sensor which helps to automatically adjust the screen brightness to a comfortable level. The data that these sensors provide pose no threat when used for their intended purpose.

With the advent of crowd sensing, this data is collected indiscriminately in order to find trends or discover interesting correlations in the data and are often kept in large databases where malicious entities can use it for nefarious purposes by revealing the identity of the persons who generated this data. For this reason, there has been a noticeable effort in the research community to develop methods and strategies to protect the privacy of the users while still being able to collect usable data from them. These methods can introduce limitations in the utility of the data and in, some cases, a non-negligible overhead in the overall data collection and data mining processes, therefore, it is advantageous to know which data has the potential to be a threat to a user's privacy so that only that data and no other is treated with the privacy-preserving methods that have been developed.

In this paper, we seek to identify what types of sensor data can be collected on a smartphone and which of those types can pose a threat to user privacy. We identify the data types by first

looking at the hardware specifications of a typical smartphone and then looking at the Android API to see what information can be retrieved from this hardware. To determine the threat level of each type we look into the literature for how this data can be used (for example, in behavioural biometrics, inference attacks, behaviour modeling, etc.). Considering the large scope of the topic we choose to focus on novel or recent work which is based on or improves upon older work.

2 Smartphone Data Types and Privacy

After analyzing hardware information of popular smartphones by Samsung, Google, LG, HTC, and Huawei we present in this section the most common sensors and data available. The information is presented in no particular order.

For certain data types we equate their potential for a privacy breach with their usefulness as authentication modalities. Our reasoning is that if a particular data type is unique enough to be used for authentication purposes, then it is certainly unique enough to identify a person in a large dataset. Furthermore, if a particular data type is significantly correlated with another data type which was evaluated to be a privacy threat, then we conclude that this particular data type must also be a privacy threat to some level.

Sampling Types For each data type we also note the type of sampling that is required to derive useful information. In table 1 we define several sampling types that can be used to collect information from sensors. We can use these sampling types to qualify the threat-sensitivity of a particular sensor’s data. For example, if some sensor only requires Type A sampling, where a single sample is enough to derive some feature, then we can conclude that it a sensor which can easily compromise an individual’s privacy and should avoid sharing any of its data or be extremely selective over which data to share and to whom. On the other hand, if a sensor requires Type F sampling, several samples can be shared without compromising the privacy of an individual.

Type ID	Sampling Type Description
Type A	Single sample is enough to derive feature confidently
Type B	Single sample is enough to derive feature somewhat confidently while several samples can improve confidence
Type C	Sampling only if there is a significant change is enough to derive feature confidently
Type D	Sampling at regular intervals during a day are enough to derive feature confidently
Type E	Continuous sampling for the duration of the action is enough to derive feature confidently
Type F	Continuous sampling without limitations is needed to derive feature confidently

Table 1: Sampling types.

2.1 GPS and location

The Android API package *android.location* is a comprehensive package that integrates GPS, WiFi APs, and Cell Identity information in a proprietary method to provide an accurate location estimate. The *Location* class in this package exposes methods to provide longitude with *getLongitude()*, latitude with *getLatitude()*, the accuracy of the estimate with *getAccuracy()*, the altitude with *getAltitude()*, the bearing with *getBearing()*, and even the speed with *getSpeed()*.

Location data has received a significant amount of attention from the research community in the context of privacy. It is seen as a major threat to individual privacy and at the same time its utility is undeniable as evidenced by the vast number of location-based services available. Krumm[1] has outlined some of the threats posed by location data. Someone can infer significant places like home and work, and more recently, Do et al.[2] were able to reliably characterize 10 categories of places of a person's everyday life, these included home and work as well as friend's home, transportation, friend's work, outdoor sport, indoor sport, restaurant or bar, shopping, holiday. Krumm shows examples of how pseudonymized or anonymized location data can still be used to identify the people in the data. Other information such as mode of transportation (bus, foot, car, etc.), age, work role, work frequency, and even smoking habits can also be inferred from location data. The evidence for the privacy risks of location data is overwhelming. For the GPS sensor, Type A sampling is enough to reveal the location while type C and D is enough for detecting personal and significant places.

2.2 Telephony

The network antenna is used to connect to the cellular network (GSM, edge, HSPA, LTE, etc.). The Android API *android.telephony* package can be used to get information such as the identity of the cell tower which the phone is connected to (Cell ID) and the signal strength to this cell using the method *getAllCellInfo()* from the *TelephonyManager* class. This class can also provide the service state with *getServiceState()*, network type with *getNetworkType*, call state with *getCallState()*, and data state with *getDataState()*.

The Cell ID can be used in conjunction with publicly available data of their locations to localize a person as demonstrated by LaMarca et al.[3]. Although a single sample is usually enough to determine an approximate location, several samples might be needed to increase confidence (Type B sampling). As such, the same privacy threats as location can be applied here. However, even without knowing the location of the Cell IDs, one can infer places such as home and work as done by Yadav et al.[4] as well as our own work [5]. Furthermore, since a person's connection traces to Cell IDs is directly related to the person's location traces, the Cell ID traces can be thought of as a quasi identifier much like location. For this, Type C or D sampling is required.

2.3 Bluetooth

The bluetooth antenna is used to connect to nearby bluetooth devices such as wireless headphones or a smartwatch. The Android API *android.bluetooth.le* package can be used to get a list of nearby bluetooth devices using the *startScan(...)* method of the *BluetoothLeScanner* class. This method returns a list of class *Scan-Result* which includes the hardware ID of the bluetooth devices with the *getDevice()* method and the signal strength with the *getRssi()* method.

Bluetooth connections to personal devices such as headphones and smartwatch are, in general, unique to each individual, as such, they can be used as identifying information. Bluetooth devices in range (not necessarily connected to) are not as unique but provided that some of those Bluetooth devices are geographically stationary then a frequent Bluetooth device scan can also be used to crudely localize a person as demonstrated again by LaMarca et al.[3]. Type B sampling is recommended for localization, while Type C or D is required to detect personal or significant places.

2.4 WiFi Antenna

The WiFi antenna is used to connect to WiFi networks. The Android API *android.net.wifi* package can be used to get a list of WiFi access points (APs) using the *startScan()* method from the *WifiManager* class. This method returns a list of class *ScanResult* which include the AP identity in the *BSSID* public field and the signal strength in the *level* public field.

WiFi connections to personal access points (APs) such as someone's home or work, much like Bluetooth, can be unique for each individual. It has also been demonstrated that WiFi APs in range and their signal strength can be used to localize a person by LaMarca et al.[3] and Redzic et al.[6] among many others [7, 8]. The same sampling requirements as Bluetooth apply for WiFi.

2.5 Touchscreen

The touchscreen is the main input method on a smartphone, it is used to select items on the screen, to type text, or other gestures which are out of the scope of this work. The Android API package *android.view* includes the class *View.OnTouchListener* which can be used to capture touch events. For security reasons the location of the touch is only available to the application on the foreground, but the touch event itself can still be useful information.

The dynamics of touch events (time between touches, duration of touch, pressure, etc.) are categorized as *keystroke dynamics* and they have been researched heavily for authentication and user recognition for hardware computer keyboards and more recently for smartphones [9, 10, 11]. Frank et al.[12] show that touchscreen data like navigational strokes (a subset of keystroke dynamics since they do not include typing) cannot be reliably used for authentication as a standalone but provides useful authentication features nonetheless and using this kind of data for authentication is ultimately feasible. Antal et al.[13] and Roh et al.[14] among others[15] have shown that keystroke dynamics along with additional features that can be collected on a smartphone (accelerometer, pressure, finger area) can be used to improve the performance of authentication. Continuous sampling for the duration of the keystrokes is required for their detection (Type E sampling).

2.6 Microphone

The microphone is used to capture audio to facilitate a phone call or to record audio. The Android API *android.media* package includes the class *AudioRecord* which can be used to capture the audio from the microphone. For the below mentioned exploitation methods of audio signals, continuous sampling is required for the duration of the action in order to apply the methods (Type E sampling).

The audio of someone speaking can be used to recognize them. Speaker recognition is a well researched topic, low level features like short-term spectrum and mel-frequency cepstral coefficients, voice source feature estimation, formant transitions, prosodic features, and high level features such as lexicon have been used in models like vector quantization (VQ), Gaussian mixture models (GMM), support vector machines (SVM), and neural networks [16]. More recently, with the advent of deep learning, more complex and robust modeling techniques have emerged [17, 18]. Speaker recognition has reached a high enough technological maturity level that it has found commercial applications in automated home assistants such as the Google Home.

Many human activities produce characteristic sounds which can be used to recognize them. Activities such as cooking, brushing teeth, showering, washing hands, urinating, shaving, drinking, etc. have been shown to be recognizable by the sounds they produce by several researchers [19, 20, 21]. More impressively, not only can someone recognize the activity of typing on a

physical keyboard but also recognize what is being typed solely from the data of a microphone [22, 23, 24].

Environmental noise features from audio recordings can be used to identify the location of the recording. Acoustic environment identification (AEI), as it is commonly known, is mostly limited to room or enclosed space environments where the geometry of the room can have noticeable effects on the reverberation of the audio. The main applications of AEI are in audio forensics where an estimation of the reverberation and background noise from a recording can be used to identify the room or even the location inside a room where the audio was recorded [25, 26, 27, 28]. Prior measurements or estimates of the impulse response of the rooms are required for these methods since they describe how the sound reverberates in that room.

Since the room geometry can affect the audio reverberation patterns of a room, someone could use an audio recording of a sharp noise (like a hand clap) to estimate the impulse response of a room and then estimate the dimensions or even the shape of the room [29, 30, 31]. These methods are often tested under controlled environments and with specialized audio equipment so it is unclear whether a recording from a smartphone microphone would be sufficient for meaningful results.

2.7 Camera

There are often two cameras on a smartphone, the front facing camera and the main camera on the backside of the phone. They are used to take pictures, video, and to facilitate video calls. The Android API `android.hardware.camera2` package provides the necessary methods to retrieve data from the camera.

Pictures or video from a camera can be used in a several different ways to reveal information about the user even without the use of the file metadata. The most obvious is if the subject of the picture is the user themselves or of people related, in the social sense, to the user. If the subject of the picture is a city, a street, or a landmark, algorithms can be used to match the pictures to a location provided there is a database of prior pictures in that location [32, 33, 34]. There are also algorithms that can recognize the style of an image and match it to a known photographer [35, 36]. Since a single picture is used in these cases, Type A sampling is enough. Videos can also be used with the aforementioned techniques by treating them as sequences of still images. In addition, analyzing the device movement from a video can also be used to identify the user similar to gait recognition in other behavioural biometric identification schemes [37]. Type E sampling is required for this.

2.8 Environmental and Activity Sensors

There is a variety of environment and activity sensors on smartphones. Their data is exposed in the Android API `android.hardware` package with the classes `SensorManager`, `Sensor`, and `SensorEvent`. Each sensor type is assigned an integer identifier constant with an appropriate name. Among these sensors are *software sensors*, that is, sensors that do not have a direct hardware counterpart but are calculated from the outputs of one or more hardware sensors. These sensors do not require any special permissions to be accessed which makes it easy for a rogue application or website to get this data without the user's knowledge.

Many of these sensors are based on microelectromechanical systems (MEMS) technology which has been shown to be vulnerable to sensor fingerprinting [38, 39, 40, 41, 42]. The accelerometer, gyroscope, magnetometer, and barometer are all based on MEMS technology. The idea behind sensor fingerprinting is that minor manufacturing defects give each sensor a unique output which is composed of the true reading (acceleration, magnetic field strength, etc.) plus

the bias caused by the manufacturing defect. This makes it so that someone can discriminate the devices which produce a given sensor output. To achieve this, Type E sampling is required. In the sections below we will take a look at each individual sensor for their respective privacy threats which are additional to the aforementioned sensor fingerprinting.

Accelerometer The accelerometer (TYPE_LINEAR_ACCELERATION) is a hardware sensor that measures linear acceleration. It's main uses include adjusting the display orientation to match the orientation of the physical display and as a step counter among others. To derive other more interesting information besides the orientation of the device, Type E sampling would be required. The accelerometer can be used in a variety of ways to become a threat to one's privacy. It has found uses in indoor localization systems where GPS is not available. Together with gyroscope and/or magnetometer readings it can help to accurately track the movement of a person [43, 44, 6]. It is often used for activity recognition as well (sitting, walking, running, biking, cleaning, shopping, sleeping, cooking, etc.) [45, 46, 47, 48]. Its applications also extend into behavioural biometrics where gait recognition uses the accelerometer to recognize a person based on how they walk or move [49, 10, 11]. When coupled with touch event detection it has even been used to detect what is being typed on the touch screen [50]. Therefore the accelerometer can reveal not only location, but activity patterns throughout one's daily life, the identity of someone based on how they walk, and in some cases, even what they type on their smartphone. It has been shown that auditory vibrations can be picked up by the accelerometer on modern smartphones like the iPhone 4 or a Samsung Galaxy S4 and can be used to detect hotwords (short keywords or phrases that are often used to activate voice assistants) or even what is being typed on a physical keyboard nearby [51, 52].

Gyroscope The gyroscope (TYPE_GYROSCOPE) is a hardware sensor that measures the rotation or twist of the device. It is often used in conjunction with the accelerometer to measure the orientation of the device and to aid in navigation/localization schemes. Michalevsky et al.[53] show that sounds can affect the measurements of a gyroscope to such a level that private information about the phone's environment can be revealed such as who is speaking and to some extent, what is being said. Type E sampling is required for these methods.

Magnetometer The magnetometer (TYPE_MAGNETIC_FIELD) is a hardware sensor that is mainly used to measure the Earth's magnetic field for the purpose of navigation. It has found uses in indoor localization schemes by comparing the magnetic field to previously collected magnetic field fingerprints to localize a person [54, 55, 56]. These methods require Type E sampling and prior data collection to map the fingerprint to specific locations. It is not applicable for outdoor environments since these methods rely on the structural supports of building and rooms which produce these magnetic fingerprints. For outdoor environments it can only reliably measure the orientation of the smartphone with respect to the Earth's magnetic field.

Barometer The barometer (TYPE_PRESSURE) is a hardware sensor that measures the atmospheric pressure. Not all devices are equipped with this sensor. Barometric pressure varies depending on the weather and on altitude. Baring extreme weather events, the rate of change of barometric pressure due to weather is relatively slow (less than 0.04hPa per hour for steady weather, less than 0.5hPa per hour for slow weather changes, and up to 3hPa per hour for rapid weather changes). While in a city like Geneva, Switzerland where the highest altitude is 457m and lowest is 370m, one can expect a change of approximately 0.115hPa per meter of altitude change. Based on these crude estimates it is no surprise that the barometric pressure is often

used as an altimeter and with its inclusion in smartphones it has aided in indoor navigation algorithms to determine the floor that the person is on [43, 57, 58, 59]. As such, someone with access to barometer data can learn about the altitude or floor in which a person lives and works as well as altitude variations during their commute. The specific methods vary in their sampling from Type B to Type F. For a city with many altitude variations like Geneva, it does not seem out of the realm of possibility to be able to reconstruct the commute path of a person based on barometric data, it is something worth looking into.

Proximity The proximity sensor (TYPE_PROXIMITY) is a hardware sensor that measures distance. It is mainly used to detect when the user places the device next to their ear during a phone call so that the screen can be turned off in order to save power. In most cases the sensor has a very limited range of up to 5cm and only tells you if there is something near it (less than 5cm). As such, it is only useful to know if the phone is in a pocket, bag, or next to your ear when taking a call. It does not appear to have any immediate implications to privacy.

Ambient light The ambient light sensor (TYPE_LIGHT) is a hardware sensor that measures the intensity of light. It is mainly used to automatically adjust the screen brightness to a comfortable level. Ambient light during daytime varies significantly for indoor and outdoor locales, therefore, someone can easily detect this during the daytime using this sensor [60]. Type C or D sampling would be enough to detect when the user changes from indoor to outdoor throughout the day. Kayacik[61] and Micallef et al.[62] created temporal and spatial models for light sensor readings among other sensors and their results show that the light sensor readings are among the sensors with the highest similarity between users. Based on their results they conclude that, on its own, the light sensor is not sufficient for authentication. An interesting exploit of the ambient light sensor was revealed by Spreitzer[63] where they showed that by using variations in the ambient light due to slight tilting of the smartphone while inputting a PIN they can improve their chances of correctly guessing it. They used a corpus of 50 random PINs and allowed themselves 10 guesses and managed to have an 80% success rate compared to 20% if they randomly guessed. Type E sampling during the PIN entry was used. The ambient light sensor has also found a use in indoor localization. If one has control of the LED lighting in a room they can send detectable light variations to the phone and help it to localize itself in the room [64]. Mazilu et al.[65] have also shown that it is feasible to detect room changes solely based on the ambient light sensor readings. Both of these indoor localization methods require Type E sampling.

Gravity The gravity sensor (TYPE_GRAVITY) is a software sensor that provides the direction and acceleration due to gravity. It most commonly uses the readings of the accelerometer and the gyroscope. It is directly correlated with the physical orientation of the device. The main use of this software sensor is to remove the gravity component from raw accelerometer measurements and be able to use those measurements for other tasks that require only the linear acceleration. On their own, the gravity measurements have very little utility and therefore do not pose any apparent threat to privacy.

Step The step sensor (TYPE_STEP_COUNTER, TYPE_STEP_DETECTOR) is a software sensor that detects when the steps a user makes when walking. It uses the accelerometer readings to derive the steps. When stride length is known (distance after one step) or accurately estimated from the height of a person, step counts can be used to estimate the distance that a person has walked [66, 67, 68, 69]. Since only one sample is needed to derive the distance, Type A sampling

is enough. Although there is significant error depending on what device is being used or even depending on the speed that a person is walking, someone can roughly determine the distances to nearby destinations where the user walks to. There are no significant privacy concerns for this data since the accuracy of these measurements can have significant errors over longer distances or even at different walking speeds.

2.9 Summary

In table 2 we summarize the possible threats of each sensor noting the type of sampling that is required. Location and location features seem to be a common type of threat for most sensors. In table 3 we summarize the literature which was used.

Sensor	Threat Summary	Sampling Reqs
GPS	location and personal places[1, 2]	Type A for location, Type C and D for personal places
Cell ID	location and personal places[3, 70, 5]	Type B for location, Type C and D for personal places
Bluetooth	location and personal places[3], identity (from connections to personal devices)	Type B for location, Type C and D for personal places and identity
WiFi	location and personal places[3, 7, 6, 8], identity (from connections to personal devices)	Type B for location, Type C and D for personal places and identity
Touchscreen	identity (keystroke dynamics[9, 12, 10, 13, 11, 14, 15])	Type E
Microphone	identity (speaker recognition[16, 17, 18]), activity[19, 20, 21], keylogger (for physical keyboard[22, 23, 24]), location features (AEI[25, 26, 27, 28], room characteristics[29, 30, 31])	Type E
Camera	location and location features[32, 33, 34], identity (selfies, gait recognition from video[37], author recognition[35, 36])	Type A for static pictures, Type E for video
All MEMS	identity (MEMS sensor fingerprinting[38, 39, 40, 41, 42])	Type E
Accelerometer (MEMS)	location (indoor navigation[43, 44, 6]), activity[45, 46, 47, 48], PIN[50], identity (gait recognition[49, 10, 11], speaker recognition[51, 52])	Type E
Gyroscope (MEMS)	identity (speaker recognition[53])	Type E
Magnetometer (MEMS)	location (indoor localization via fingerprinting[54, 55, 56])	Type E
Barometer (MEMS)	location features (floor detection[43, 57, 58, 59])	Type B up to Type F
Proximity	None	
Ambient light	location features (indoor vs outdoor[60], indoor navigation[64], room detection[65]), PIN [63]	Type C and D for indoor/outdoor/room features, Type E for navigation and PIN
Gravity	None	
Step	distance walked (estimated from number of steps [66, 67, 68, 69])	Type A

Table 2: Summary of sensors and corresponding privacy threats

Citation	Sensors used	Derived information
[2] Details	GPS, WiFi, Bluetooth, App	Location of home, work, other personal places Ground truth: User annotated data. Methodology: Random forest classifier. Accuracy & Limitations: GPS features alone gave 70.3% accuracy, adding Wifi features to previous 71.7%, adding Bluetooth features to previous 74.6%, adding app features to previous 75%. Infrequently visited places are not reliably recognized.
[3] Details	WiFi, Bluetooth, Cell ID	Map of radio beacons, location of user Ground truth: GPS war-driving or institution databases with location of radio beacons. Methodology: tracker component that models signal propagation and takes into account physical environment (for example, buildings). A probabilistic Bayesian particle filter can be used to increase accuracy. Accuracy & Limitations: lower accuracy than GPS.
[70] Details	GPS, WiFi, Bluetooth, Cell ID	Location personal places Ground truth: user annotated data. Methodology: GPS or PlaceLab estimated location was used to collect traces. Time based clustering was used on location traces to find personal places. Accuracy & Limitations: Does not label the personal places.
[5] Details	Cell ID	Detection of personal place Ground truth: GPS and user annotated data. Methodology: graph based clustering of Cell IDs using Cell ID transition matrix populated by Cell ID oscillation events. Duration of stay in clusters and time of day indicating home or work. Accuracy & Limitations: limited to urban environment with relatively dense cellular tower deployment. Does not detect places with shorter durations of stay.
[6] Details	WiFi	indoor location Ground truth: ground truth. Methodology: Fingerprinting of RSSI of WiFi access points at specific calibration points (CPs) and using naive Bayes to identify the three nearest CP, then using interpolation driven by the likelihoods to find the location of the user in the vicinity of those CPs (even using as few as 2 of them). Accuracy & Limitations: Accuracy is around 2 meters which can be significant in indoor environments even though they showed that this method is better than many others. Requires calibration measurements in advance.
[12] Details	touchscreen (navigational strokes)	user identity Ground truth: 41 users read text and compare images on an android phone to produce natural navigational strokes. Methodology: 30 behavioural touch features (for example, mid-stroke area covered, direction of end to end line, start/end x, start/end y, and more). From sets of highly correlated features, only one was selected. Used kNN and SVM classifiers. Accuracy & Limitations: 0% to 4% error (false negative and false positive combined) which is not ideal for authentication purposes. More subjects needed to improve feature selection. Differences of screen sizes of devices needs to be taken into account.
[13] Details	touchscreen	user identity Ground truth: 42 users. Android application with its own keyboard. Nexus 7 tablet (37 users) and LG Optimus L7 II p710 phone (5 users). Users input a password 30 times (same for all). Methodology: features: time between key press and release, time between consecutive key presses, time between key release and next press, pressure of press, finger area of press, averages of previous values. WEKA machine learning software was used. Analyzed several classifiers. Accuracy & Limitations: Best classifier was random forest with 82.53% accuracy using only time based features, and 93.04% accuracy using time based features and touchscreen based features together.

Continued on Next Page. . .

Citation	Sensors used	Derived information
[15]	touchscreen, accelerometer, gyro-scope, magnetometer	user identity
Details	<p>Ground truth: 100 users typing 3 answers of at least 250 words under sitting or walking conditions. Sensor sampling at 100Hz.</p> <p>Methodology: Scaled Manhattan (SM), scaled Euclidean (SE), SVM verifiers using hand movement, orientation, grasp (HMOG) features, tap and keystroke dynamics features.</p> <p>Accuracy & Limitations: Best verifier was SM with Equal Error Rate of 10.05% for sitting and 7.16% for walking. Including HMOG features improved accuracy over only tap or keystroke dynamics. Cross-device interoperability and varying walking speeds were not explored.</p>	
[20]	microphone	Human activity (cleaning, brush teeth, walk, drink water, etc.)
Details	<p>Ground truth: Sound recordings of each activity</p> <p>Methodology: 5 random segments of 1.5 second from recording were used. Mel Frequency Cepstral Coefficients (MFCC) were extracted for each segment. Discrete time warping was used to get closest match.</p> <p>Accuracy & Limitations: Average accuracy of recognizing each of the 14 activities was 92.5% (80% lowest, 100% highest). Sound samples were recorded in a controlled environment, realistic data would improve argument.</p>	
[24]	microphone	text typed on physical keyboard
Details	<p>Ground truth: 10 minute recording of user typing in English</p> <p>Methodology: Compute Cepstrum features of each keystroke. For training, use clustering technique to separate into classes and language model correction based on HMM to label and then train a classifier. For recognition, use classifier and language model correction.</p> <p>Accuracy & Limitations: 90% of 5-character passwords in fewer than 20 attempts, 80% of 10-character passwords in fewer than 75 attempts. Classifiers user: linear classification, Gaussian mixtures, or Neural Network.</p>	
[27]	microphone	environment (room)
Details	<p>Ground truth: 30 audio recordings in 6 different acoustic environments (big classroom 1 and 2, small classroom, small seminar hall, seminar hall, small room)</p> <p>Methodology: Blind de-reverberation was used to extract reverberant component of audio. Impulse response was estimated via hand-clap method. MFCCs were used as features, a multiclass SVM was used for classification.</p> <p>Accuracy & Limitations: 4 rooms identified with 100% accuracy, 2 rooms above 80% accuracy. Need to measure impulse response of rooms separately. Environments were based only on university campus.</p>	
[30]	microphone	room dimensions
Details	<p>Ground truth: simulations of rectangular and L-shaped room.</p> <p>Methodology: Defined a cost function robust against wrong matches of TOAs. Genetic algorithm was used to minimize cost function and derive room dimensions.</p> <p>Accuracy & Limitations: room dimensions for rectangular room are within 10cm of actual size, for L-shaped room within 70cm. Should be repeated in real room. Room shape known a priori.</p>	
[34]	camera (photo)	location
Details	<p>Ground truth: 126M photos with Exif geolocations from the web.</p> <p>Methodology: Used convolutional neural network (CNN) to train with 91M images, the rest used for validation. 237 geotagged Flickr photos used to measure accuracy of model.</p> <p>Accuracy & Limitations: When using any type of photo accuracy is 8.5% for 1km radius, 24.5% for 25km, 37.6% for 200km, 53.6% for 750km, 71.3% for 2500km. Using other contextual info increased accuracy.</p>	

Continued on Next Page...

Citation	Sensors used	Derived information
[37]	camera (video)	identity
Details	<p>Ground truth: 32 users recorded two 7 minute sequences with head-mounted cameras.</p> <p>Methodology: optical flow vectors computed for each frame. CNN with 2 hidden layers for classifier.</p> <p>Accuracy & Limitations: 77% accuracy for 4 second of video, 90% accuracy for 12 seconds of video. Stabilizing the video deteriorated results. Requires that camera be mounted on person. Should consider hand-held camera.</p>	
[39, 41, 42]	MEMS (accelerometer, gyroscope, magnetometer)	device identity
Details	<p>Ground truth: 3 devices on a robotic arm and moved in a predetermined pattern. For magnetometer, 9 devices were tested, a solenoid was placed around each device and a predetermined signal was produced.</p> <p>Methodology: SVM classifier was used with different kernel functions.</p> <p>Accuracy & Limitations: Over 95% accuracy to distinguish between different models, over 65% accuracy overall. The inputs to the sensors were controlled. This might not be possible to apply with data collected in the wild.</p>	
[46]	accelerometer	human activity (walking, jogging, ascending stairs, descending stairs, sitting, standing)
Details	<p>Ground truth: 29 users performing each activity several times while carrying a smartphone.</p> <p>Methodology: Split data into 10 second segments, each segment extracted features like average acceleration, standard deviation, time between peaks, etc. WEKA with decision trees (J48), logistic regression, multilayer neural networks (NN) with default settings.</p> <p>Accuracy & Limitations: NN is best with an average of 91.7% accuracy. Up/down stairs had the worst accuracy as low as 44.3% and were most often confused with each other or walking. Activity set is limited, different carrying patterns of device not taken into account.</p>	
[50]	accelerometer, gyroscope	smartphone keyboard input
Details	<p>Ground truth: 10 users using a custom application for tapping icons and typing text (each letter 50 times, 19 different pangrams, and 20 times the same pangram).</p> <p>Methodology: Detect taps and extracts features of each tap (time domain and frequency domain). kNN, multinomial logistic regression, SVM, random forests, bagged decision trees are all used together in an ensemble classifier.</p> <p>Accuracy & Limitations: 90% accuracy for inferring tap locations, 80% accuracy for letters. The classifier is resource heavy and could have redundancies.</p>	
[51]	accelerometer	text typed on physical keyboard
Details	<p>Ground truth: iPhone placed on same surface as keyboard. Sentences typed were selected from the Harvard Sentences corpus.</p> <p>Methodology: Features from keypress data were used like root mean square, skewness, variance, kurtosis, FFT, MFCCs. Two neural networks were trained with with a difference in features used.</p> <p>Accuracy & Limitations: Tested with and without dictionary knowledge and with a news article from a newspaper. As much as 80% accuracy of typed content with the use of dictionary. Orientation of device, desk surface material, typing speed, ambient vibrations can affect the performance.</p>	

Continued on Next Page...

Citation	Sensors used	Derived information
[52]	accelerometer	hotword detection (for example, "okay Google")
Details	<p>Ground truth: 10 users recorded saying "Okay google" and 20 common short phrases 10 times each. Each recording played through phone speakers 10 times for training at 70dB.</p> <p>Methodology: 2 second window is used and time domain and frequency domain features are extracted. For mobile scenario, a high pass filter with 2Hz cutoff is used to remove effects due to movement. A decision tree classifier is used.</p> <p>Accuracy & Limitations: 85% in static scenario, 80% in mobile scenario. The mobile scenario is very limited with just a controlled walking. More complicated movements make it vastly more difficult to recognize the hotwords.</p>	
[53]	gyroscope	identity, speech
Details	<p>Ground truth: Nexus 4, Nexus 7, Samsung Galaxy S III were used. A loud-speaker at 75dB. TIDIGITS corpus was used (recordings of 10 users speaking the 11 digits twice each).</p> <p>Methodology: 10-30ms sliding windows with time domain features and MFCCs and Short Time Fourier Transform (STFT). SVM, GMM, DTW were used as classifiers.</p> <p>Accuracy & Limitations: Over 80% accuracy for gender ID using SVM. Speaker ID ~50% accuracy using DTW with Nexus 4 but 17% with Samsung. Speaker-independent word rec performed poorly, but improved to 65% using DTW with speaker-specific models. Results varied significantly between devices.</p>	
[55]	magnetometer	location (indoor)
Details	<p>Ground truth: Magnetic field map data collected by following serpentine pattern in a room on x-axis and then on y-axis. Test data collected following a well defined straight line, or circle path.</p> <p>Methodology: Magnetic field map data was used to generate a map of the field in the room. Test data was then matched to the map using a particle filter.</p> <p>Accuracy & Limitations: Within 0.7m of ground truth. Wi-Fi was used to get coarse location as initial condition for particle filter.</p>	
[56]	magnetometer	location (indoor)
Details	<p>Ground truth: 2 users with HTC Nexus One phone. Magnetic fingerprints collected in hallways of campus buildings as users walked along the walls and pillars.</p> <p>Methodology: Magnetic field fingerprints were collected and then DTW was used on test data to match to the fingerprints.</p> <p>Accuracy & Limitations: Hallways were detected with over 90% accuracy after only less than 5 meters of walking. Users were instructed to walk close to objects that influence magnetic fields like pillars.</p>	
[59]	barometer	location (floor in building)
Details	<p>Ground truth: 63 trials at 5 different tall buildings in New York City where barometric pressure was recorded and random floors were selected. The user could choose either the staircase or elevator.</p> <p>Methodology: Calculated the change in height based on the international pressure equation. To resolve to a floor number they used calculated clusters derived from data of all visits to building (floor height could be estimated).</p> <p>Accuracy & Limitations: 65% accuracy when floor height is not known and a default 4.02m was used (98% within 1 floor), 100% accuracy if floor height has been previously estimated.</p>	

Continued on Next Page...

Citation	Sensors used	Derived information
[65]	ambient light sensor	location (indoor room detection)
Details	<p>Ground truth: 3 users with Samsung Galaxy S4 collected data in their homes. Users logged room label each time they entered a new room on paper-based diary. Total of 132 hours of data</p> <p>Methodology: If sensor data feature was higher than a fixed threshold then a room change was detected. Decision trees (C4.5) were used for room identification.</p> <p>Accuracy & Limitations: Using only light sensor, accuracy was around 50%, with additional sensors like temperature and humidity the accuracy was above 60%. Random guess was 25% accuracy at best. Time of day, weather, and open windows affected the performance.</p>	
[63]	ambient light sensor	PIN
Details	<p>Ground truth: Samsung Galaxy SIII was used. 29 test runs by 10 users who entered 15, 30, or all 50 of the random PINs from 3 to 10 times.</p> <p>Methodology: Multiclass logistic regression, discriminant analysis, and K-nearest neighbor methods were used on the collected data with only light intensity and with additional RGBW information which modern light sensors include.</p> <p>Accuracy & Limitations: 80% success after 10 guesses from a set of 50 PINs. The set of 50 PINs is unrealistic as there many more possible combinations.</p>	

Table 3: Summary of selected state of the art

3 Discussion

After reviewing each data type in section 2 we conclude that most of them can be used on their own to reveal something about a user be it small (for example, the floor on a building) or big (for example, the location of their home and work). Combining different data types can enhance the precision, or accuracy, or both as evidenced by several of the surveyed research in table 3.

On the Android OS there is a permission framework to enable an application to explicitly request from the user if a certain data type can be used or not. Permissions that have a protection level of *normal* are automatically granted by the system while those that have protection level *dangerous* require the user’s explicit permission to be allowed. In the older versions of the Android OS, the permissions were requested in batch when installing an application but on the latest Android OS version the permissions are requested individually and on an as-needed basis during the runtime of the application (i.e. until the application needs to use the microphone it will not ask for permission). Furthermore, on the latest Android OS version, a user can adjust the individual data type permissions in the settings for each application after the fact. Consequently, the user is informed about the various data types that an application uses. The permission framework does not cover the sensors in section 2.8 at this moment and it is unclear if it will in the future.

Christin et al.[71] summarize countermeasures to several privacy threats: tailored sensing and user preferences, pseudonymity, spatial cloaking, hiding sensitive locations, data perturbation, data aggregation, among others. They also present important research challenges in this field that have yet to be fully addressed.

This and most such privacy research are concerned with threats in the context of data collection campaigns for research and data mining but similar principles must also be applied to commercially available smartphone applications. Each year Google has to remove more and more malicious applications from their marketplace amounting to hundreds of thousands of applications[72, 73]. Although a lot of malicious applications are automatically filtered, some

Data type	permission	prot. level	comments
Location (precise)	ACCESS_FINE_LOCATION	dangerous	
Location (approximate)	ACCESS_COARSE_LOCATION	dangerous	
Network Cell ID	ACCESS_COARSE_LOCATION	dangerous	
Bluetooth APs	BLUETOOTH_ADMIN	normal	
WiFi APs	ACCESS_WIFI_STATE	normal	
Touchscreen	No permissions are required	N/A	Touch event only outside application window or touch location available only to the application in foreground
Microphone	RECORD_AUDIO	dangerous	
Camera	CAMERA	dangerous	
Accelerometer	No permissions are required	N/A	
Gyroscope	No permissions are required	N/A	
Magnetometer	No permissions are required	N/A	
Barometer	No permissions are required	N/A	
Proximity	No permissions are required	N/A	
Gravity	No permissions are required	N/A	
Step	No permissions are required	N/A	

Table 4: Data types and corresponding Android OS permission requirements and protection level

may still slip through and become available for millions of people to install. Some of these can easily collect data such as location from unsuspecting users even if they have to request the specific permission from the user. So many applications require the location permission that a user might not think twice about allowing it. In table 5 we list 25 of the top installed Android applications[74] along with some of the relevant permissions that are required to fully operate them[75]; the location permission is very common.

Users should always question if an application really needs a specific permission to function. For example, location can be used for navigation, to check-in to places in social media, to show local weather, to share your location with a contact, for fitness tracking, to show location-based notifications, and many more. The issue arises when an application does not need a precise location (for example, a weather application) or only needs some tracking information (for example, a fitness application). A user should not need to give more information than is needed for an application to function much like when a stranger asks for your contact information you have the choice of giving them any of the following information depending on the intimacy level: first name, last name, email address, phone number, home address, work address, friend's address, parents' address, frequented bars, frequented shops, parents' first and last names, etc. Technically, all of these pieces of information can be considered *contact information* but you would not give out all of them if you only need to give out a first name and an email address for example. Sharing more than is necessary can feel highly intrusive. Therefore, a weather application should only get a meteorological region instead of exact coordinates, and a fitness application should only get distance and speed instead of the exact coordinates of the path you ran. There is currently no mechanism on Android OS or any other popular smartphone operating

Application	Category	Permissions
Facebook	Social	Location, Camera, Microphone, WiFi
WhatsApp	Communication	Location, Camera, Microphone, WiFi
Messenger (Facebook)	Communications	Location, Camera, Microphone, WiFi
Subway Surfers	Game Arcade	WiFi
Skype	Communication	Location, Camera, Microphone, WiFi
Clean Master	Tools	Location, Camera, Microphone, WiFi
Security Master	Tools	Location, Camera, Microphone, WiFi
Candy Crash Saga	Game Casual	WiFi
UC Browser	Communication	Location, Camera, Microphone, WiFi
Snapchat	Social	Location, Camera, Microphone, WiFi
My Talking Tom	Game Casual	Microphone, WiFi
Twitter	News & Magazines	Location, Camera, Microphone, WiFi
Viber Messenger	Communication	Location, Camera, Microphone, WiFi
LINE	Communication	Location, Camera, Microphone, WiFi
Pou	Game Casual	Microphone, WiFi
Super-Bright LED Flashlight	Productivity	Camera
Temple Run 2	Game Action	WiFi
SHAREit	Tools	Location, Camera, WiFi
imo free video calls and chat	Communication	Location, Camera, Microphone, WiFi
Microsoft Word	Productivity	Camera, WiFi
Flipboard: News For Our Time	News & Magazines	
Clash of Clans	Game Strategy	WiFi
Spotify Music	Music & Audio	Camera, Microphone, WiFi
Shadow Fight 2	Game Action	WiFi
Pokemon GO	Game Adventure	Location, Camera

Table 5: Top downloaded apps excluding pre-installed and system applications. Only the following permissions were noted: Location, Camera, Microphone, and WiFi.

system that provides this level of abstraction when it comes to location information, location context information, or most other types of data that can be collected on a smartphone device.

4 Implications for GDPR Compliance

The GDPR applies to any entity that handles, uses, or collects "personal data". In the GDPR, personal data is loosely defined as "*any information relating to an identified or identifiable natural person ('data subject'); an identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person*". This definition is general enough to cover a wide range of data but for non-experts it is not immediately obvious which data is **actually** covered.

Each EU member state is responsible for enforcing the GDPR rules by appointing a supervisory authority (SA) who works with other member state SAs to keep consistency among them. The European Data Protection Board coordinates the SAs. Individuals can submit GDPR claims to the relevant SA who will evaluate the claim and proceed with the appropriate actions. The SA also provides some basic guidance to businesses and organizations in order to help them

comply with the GDPR, for example, in the form of a self-assessment checklist. Still, no specific definition is given for what constitutes personal data.

Regarding this issue, in this work, we sought to bring some clarity for smartphone sensor data, a rather small subset of all the possible data types out there. We saw that most of the sensor data on a smartphone can reveal personal information but we can get a sense of how sensitive they are by using the sampling types that we defined (see table 1) and assigned to each data type. Data that require a sampling type A being very privacy-sensitive and data that require a sampling type F being less privacy sensitive. The main implication of the sampling type categorization is that if an organization collects a certain data type at a lesser sampling type than is required to derive sensitive information (see table 2), then they may not have to consider this data as personal.

The GDPR has provisions for certification and certification bodies (articles 42 and 43) and a list of these can be found on the European Data Protection Board website[76]. At the time of writing of this paper there is no official GDPR certification mechanism but we can expect that there will be in the future. A GDPR certification (along with their seals and marks) can let individuals know about the GDPR compliance of an organization and reassure them that their data is handled accordingly and that they are afforded a certain control over their data. However, such a certification would not inform an individual about the level of potential privacy loss they might incur in the case of a data breach against the organization for example.

We can use the results this paper to aid in evaluating the potential privacy loss of a set of smartphone sensor data. One way to go about this is to assign numerical values to the threat posed by some type of personal information and then combine it with the sampling type required to derive this personal information. For example, lets suppose that a smartphone application only collects Cell ID data. Location and personal places can be derived from Cell ID data, so we will assume that it is a high privacy threat. The sampling type required to derive this information is C . Combining the two measures should result in a relatively high potential privacy loss. On the other hand, lets consider barometer data. This data, sampled at the sampling type B level, can be used to reveal the altitude or floor level of an individual which seems like a low privacy threat. This should result in a low overall potential privacy loss. The specifics of this evaluation can be further investigated in the future and a concrete methodology can be established for all data types, not just smartphone sensor data, and organizations can display this score in an effort to inform the users about the level of privacy loss that a user might incur in a worst case scenario.

5 Future Work

In order to have a complete summary of privacy threats from all mobile device data types, more data sources need to be investigated. In this paper we only looked at the most commonly used sensor data, but there are other sensors such as the CPU temperature and battery state which also warrant investigation. Furthermore, there are many software sources that need to be scrutinized. Some of these include application usage, phone interaction (for example, screen on/off), browsing history, instant messaging behaviour, TCP connection information, and more.

References

- [1] J. Krumm, “A survey of computational location privacy,” *Personal and Ubiquitous Computing*, vol. 13, no. 6, pp. 391–399, aug 2009. [Online]. Available: <http://link.springer.com/10.1007/s00779-008-0212-5>

- [2] T. M. T. Do and D. Gatica-Perez, “The Places of Our Lives: Visiting Patterns and Automatic Labeling from Longitudinal Smartphone Data,” *IEEE Transactions on Mobile Computing*, vol. 13, no. 3, pp. 638–648, mar 2014. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6427751><http://ieeexplore.ieee.org/document/6427751/>
- [3] A. LaMarca, Y. Chawathe, S. Consolvo, J. Hightower, I. Smith, J. Scott, T. Sohn, J. Howard, J. Hughes, F. Potter, J. Tabert, P. Powledge, G. Borriello, and B. Schilit, “Place Lab: Device Positioning Using Radio Beacons in the Wild,” in *Pervasive Computing*. Springer, 2005, pp. 116–133. [Online]. Available: http://link.springer.com/10.1007/11428572_{_}8
- [4] K. Yadav, V. Naik, A. Kumar, and P. Jassal, “PlaceMap: Discovering Human Places of Interest Using Low-Energy Location Interfaces on Mobile Phones,” in *Proceedings of the Fifth ACM Symposium on Computing for Development - ACM DEV-5 '14*, vol. 5. New York, New York, USA: ACM Press, 2014, pp. 93–102. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=2674377.2674386>
- [5] M. Fanourakis and K. Wac, “Lightweight Clustering of Cell IDs into Meaningful Neighbourhoods,” in *Performance & Security Modelling and Evaluation of Cooperative Heterogeneous Networks - HET-NETs*, D. D. Kouvatso, S. Balsamo, and Y. Takahashi, Eds. River Publishers, 2013, pp. 698–704. [Online]. Available: http://riverpublishers.com/book_{_}details.php?book_{_}id=234
- [6] M. D. Redzic, C. Brennan, and N. E. O’Connor, “SEAMLOC: Seamless Indoor Localization Based on Reduced Number of Calibration Points,” *IEEE Transactions on Mobile Computing*, vol. 13, no. 6, pp. 1326–1337, jun 2014. [Online]. Available: <http://ieeexplore.ieee.org/document/6583154/>
- [7] H. Liu, H. Darabi, P. Banerjee, and J. Liu, “Survey of Wireless Indoor Positioning Techniques and Systems,” *IEEE Transactions on Systems, Man and Cybernetics, Part C (Applications and Reviews)*, vol. 37, no. 6, pp. 1067–1080, nov 2007. [Online]. Available: <http://ieeexplore.ieee.org/document/4343996/>
- [8] Q. D. Vo and P. De, “A Survey of Fingerprint-Based Outdoor Localization,” *IEEE Communications Surveys & Tutorials*, vol. 18, no. 1, pp. 491–506, 2016. [Online]. Available: <http://ieeexplore.ieee.org/document/7131436/>
- [9] S. P. Banerjee and D. Woodard, “Biometric Authentication and Identification Using Keystroke Dynamics: A Survey,” *Journal of Pattern Recognition Research*, vol. 7, no. 1, pp. 116–139, 2012. [Online]. Available: <http://www.jprr.org/index.php/jprr/article/view/427/167>
- [10] J. Unar, W. C. Seng, and A. Abbasi, “A review of biometric technology along with trends and prospects,” *Pattern Recognition*, vol. 47, no. 8, pp. 2673–2688, aug 2014. [Online]. Available: <http://dx.doi.org/10.1016/j.patcog.2014.01.016><http://linkinghub.elsevier.com/retrieve/pii/S003132031400034X>
- [11] A. K. Jain, K. Nandakumar, and A. Ross, “50 years of biometric research: Accomplishments, challenges, and opportunities,” *Pattern Recognition Letters*, vol. 79, pp. 80–105, aug 2016. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S0167865515004365>
- [12] M. Frank, R. Biedert, E. Ma, I. Martinovic, and D. Song, “Touchalytics: On the Applicability of Touchscreen Input as a Behavioral Biometric for Continuous

- Authentication,” *IEEE Transactions on Information Forensics and Security*, vol. 8, no. 1, pp. 136–148, jan 2013. [Online]. Available: <http://ieeexplore.ieee.org/document/6331527/>
- [13] M. Antal, L. Z. Szabó, and I. László, “Keystroke Dynamics on Android Platform,” *Procedia Technology*, vol. 19, pp. 820–826, 2015. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S221201731500119X>
- [14] J.-h. Roh, S.-H. Lee, and S. Kim, “Keystroke dynamics for authentication in smartphone,” in *2016 International Conference on Information and Communication Technology Convergence (ICTC)*. IEEE, oct 2016, pp. 1155–1159. [Online]. Available: <http://ieeexplore.ieee.org/document/7763394/>
- [15] Z. Sitova, J. Sedenka, Q. Yang, G. Peng, G. Zhou, P. Gasti, and K. S. Balagani, “HMOG: New Behavioral Biometric Features for Continuous Authentication of Smartphone Users,” *IEEE Transactions on Information Forensics and Security*, vol. 11, no. 5, pp. 877–892, may 2016. [Online]. Available: <http://ieeexplore.ieee.org/document/7349202/>
- [16] T. Kinnunen and H. Li, “An overview of text-independent speaker recognition: From features to supervectors,” *Speech Communication*, vol. 52, no. 1, pp. 12–40, jan 2010. [Online]. Available: <http://dx.doi.org/10.1016/j.specom.2009.08.009http://linkinghub.elsevier.com/retrieve/pii/S0167639309001289>
- [17] Y. Lei, N. Scheffer, L. Ferrer, and M. McLaren, “A novel scheme for speaker recognition using a phonetically-aware deep neural network,” in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, may 2014, pp. 1695–1699. [Online]. Available: <http://ieeexplore.ieee.org/document/6853887/>
- [18] F. Richardson, D. Reynolds, and N. Dehak, “Deep Neural Network Approaches to Speaker and Language Recognition,” *IEEE Signal Processing Letters*, vol. 22, no. 10, pp. 1671–1675, oct 2015. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=7080838http://ieeexplore.ieee.org/document/7080838/>
- [19] Jianfeng Chen, Jianmin Zhang, A. Kam, and L. Shue, “An Automatic Acoustic Bathroom Monitoring System,” in *2005 IEEE International Symposium on Circuits and Systems*. IEEE, 2005, pp. 1750–1753. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1464946http://ieeexplore.ieee.org/document/1464946/>
- [20] Y. Zhan, J. Nishimura, and T. Kuroda, “Human Activity Recognition from Environmental Background Sounds for Wireless Sensor Networks,” *IEEJ Transactions on Electronics, Information and Systems*, vol. 130, no. 4, pp. 565–572, 2010. [Online]. Available: <http://joi.jlc.jst.go.jp/JST.JSTAGE/ieejieiss/130.565?from=CrossRef>
- [21] J. a. Stork, L. Spinello, J. Silva, and K. O. Arras, “Audio-based human activity recognition using Non-Markovian Ensemble Voting,” in *2012 IEEE RO-MAN: The 21st IEEE International Symposium on Robot and Human Interactive Communication*. IEEE, sep 2012, pp. 509–514. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6343802http://ieeexplore.ieee.org/document/6343802/>
- [22] D. Asonov and R. Agrawal, “Keyboard acoustic emanations,” in *IEEE Symposium on Security and Privacy, 2004. Proceedings. 2004*. IEEE, 2004, pp. 3–11. [Online]. Available: <http://ieeexplore.ieee.org/document/1301311/>

- [23] Y. Berger, A. Wool, and A. Yeredor, “Dictionary attacks using keyboard acoustic emanations,” in *Proceedings of the 13th ACM conference on Computer and communications security - CCS '06*. New York, New York, USA: ACM Press, 2006, p. 245. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1180436&CFID=104977675&CFTOKEN=15114097http://portal.acm.org/citation.cfm?doid=1180405.1180436>
- [24] L. Zhuang, F. Zhou, and J. D. Tygar, “Keyboard acoustic emanations revisited,” *ACM Transactions on Information and System Security*, vol. 13, no. 1, pp. 1–26, oct 2009. [Online]. Available: <http://ieeexplore.ieee.org/document/1301311/http://portal.acm.org/citation.cfm?doid=1609956.1609959>
- [25] H. Malik, “Acoustic Environment Identification and Its Applications to Audio Forensics,” *IEEE Transactions on Information Forensics and Security*, vol. 8, no. 11, pp. 1827–1837, nov 2013. [Online]. Available: <http://dblp.uni-trier.de/db/journals/tifs/tifs8.html#{#}Malik13http://ieeexplore.ieee.org/document/6595031/>
- [26] H. Zhao and H. Malik, “Audio recording location identification using acoustic environment signature,” *IEEE Transactions on Information Forensics and Security*, vol. 8, no. 11, pp. 1746–1759, 2013.
- [27] R. K. Patole, P. Rege, and P. Suryawanshi, “Acoustic environment identification using blind de-reverberation,” in *2016 International Conference on Computing, Analytics and Security Trends (CAST)*. IEEE, dec 2016, pp. 495–500. [Online]. Available: <http://ieeexplore.ieee.org/document/7915019/>
- [28] G. Delgado-Gutiérrez, F. Rodríguez-Santos, O. Jiménez-Ramírez, and R. Vázquez-Medina, “Acoustic environment identification by KullbackLeibler divergence,” *Forensic Science International*, vol. 281, pp. 134–140, dec 2017. [Online]. Available: <http://dx.doi.org/10.1016/j.forsciint.2017.10.031http://linkinghub.elsevier.com/retrieve/pii/S0379073817304334>
- [29] S. Tervo and T. Tossavainen, “3D room geometry estimation from measured impulse responses,” in *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, mar 2012, pp. 513–516. [Online]. Available: <http://ieeexplore.ieee.org/document/6287929/>
- [30] D. Markovica, F. Antonacci, A. Sarti, and S. Tubaro, “Estimation of room dimensions from a single impulse response,” in *2013 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, vol. 1, no. 1. IEEE, oct 2013, pp. 1–4. [Online]. Available: <http://ieeexplore.ieee.org/document/6701867/>
- [31] T. Rajapaksha, X. Qiu, E. Cheng, and I. Burnett, “Geometrical room geometry estimation from room impulse responses,” in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, vol. 2016-May. IEEE, mar 2016, pp. 331–335. [Online]. Available: <http://ieeexplore.ieee.org/document/7471691/>
- [32] A. R. Zamir and M. Shah, “Accurate Image Localization Based on Google Maps Street View,” in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Springer, Berlin, Heidelberg, 2010, vol. 6314 LNCS, no. PART 4, pp. 255–268. [Online]. Available: <http://link.springer.com/10.1007/978-3-642-15561-1{-}19>
- [33] G. Baatz, K. Köser, D. Chen, R. Grzeszczuk, and M. Pollefeys, “Handling Urban Location Recognition as a 2D Homothetic Problem,” in *Lecture Notes in Computer*

- Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Springer, Berlin, Heidelberg, 2010, vol. 6316 LNCS, no. PART 6, pp. 266–279. [Online]. Available: http://link.springer.com/10.1007/978-3-642-15567-3_{-}20
- [34] T. Weyand, I. Kostrikov, and J. Philbin, “PlaNet - Photo Geolocation with Convolutional Neural Networks,” in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2016, vol. 9912 LNCS, pp. 37–55. [Online]. Available: http://link.springer.com/10.1007/978-3-319-46484-8_{-}3
- [35] S. Karayev, A. Hertzmann, M. Trentacoste, H. Han, H. Winnemoeller, A. Agarwala, and T. Darrell, “Recognizing Image Style,” in *Proceedings of the British Machine Vision Conference 2014*. British Machine Vision Association, 2014, pp. 122.1–122.11. [Online]. Available: <http://arxiv.org/abs/1311.3715><http://www.bmva.org/bmvc/2014/papers/paper121/index.html>
- [36] C. Thomas and A. Kovashka, “Seeing Behind the Camera: Identifying the Authorship of a Photograph,” in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, jun 2016, pp. 3494–3502. [Online]. Available: <http://ieeexplore.ieee.org/document/7780749/>
- [37] Y. Hoshen and S. Peleg, “An Egocentric Look at Video Photographer Identity,” in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 2016-Decem. IEEE, jun 2016, pp. 4284–4292. [Online]. Available: <http://ieeexplore.ieee.org/document/7780833/>
- [38] H. Bojinov, Y. Michalevsky, G. Nakibly, and D. Boneh, “Mobile Device Identification via Sensor Fingerprinting,” aug 2014. [Online]. Available: <http://arxiv.org/abs/1408.1416>
- [39] G. Baldini, G. Steri, F. Dimc, R. Giuliani, and R. Kamnik, “Experimental Identification of Smartphones Using Fingerprints of Built-In Micro-Electro Mechanical Systems (MEMS),” *Sensors*, vol. 16, no. 6, p. 818, jun 2016. [Online]. Available: <http://www.mdpi.com/1424-8220/16/6/818>
- [40] T. Van Goethem, W. Scheepers, D. Preuveneers, and W. Joosen, “Accelerometer-Based Device Fingerprinting for Multi-factor Mobile Authentication,” in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2016, vol. 9639, pp. 106–121. [Online]. Available: http://link.springer.com/10.1007/978-3-319-30806-7_{-}7
- [41] G. Baldini, F. Dimc, R. Kamnik, G. Steri, R. Giuliani, and C. Gentile, “Identification of Mobile Phones Using the Built-In Magnetometers Stimulated by Motion Patterns,” *Sensors*, vol. 17, no. 4, p. 783, apr 2017. [Online]. Available: <http://arxiv.org/abs/1701.07676><http://www.mdpi.com/1424-8220/17/4/783>
- [42] G. Baldini, G. Steri, R. Giuliani, and V. Kyovtorov, “Mobile phone identification through the built-in magnetometers,” pp. 1–10, jan 2017. [Online]. Available: <http://arxiv.org/abs/1701.07676>
- [43] G. Lammel, J. Gutmann, L. Marti, and M. Dobler, “Indoor Navigation with MEMS sensors,” *Procedia Chemistry*, vol. 1, no. 1, pp. 532–535, sep 2009. [Online]. Available: <http://dx.doi.org/10.1016/j.proche.2009.07.133><http://linkinghub.elsevier.com/retrieve/pii/S187661960900134X>

- [44] K. Subbu, C. Zhang, J. Luo, and A. Vasilakos, “Analysis and status quo of smartphone-based indoor localization systems,” *IEEE Wireless Communications*, vol. 21, no. 4, pp. 106–112, aug 2014. [Online]. Available: <http://ieeexplore.ieee.org/document/6882302/>
- [45] A. Brajdic and R. Harle, “Walk detection and step counting on unconstrained smartphones,” in *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing - UbiComp '13*. New York, New York, USA: ACM Press, 2013, p. 225. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=2493432.2493449>
- [46] J. R. Kwapisz, G. M. Weiss, and S. A. Moore, “Activity recognition using cell phone accelerometers,” *ACM SIGKDD Explorations Newsletter*, vol. 12, no. 2, p. 74, mar 2011. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1964918http://portal.acm.org/citation.cfm?doid=1964897.1964918>
- [47] O. D. Incel, M. Kose, and C. Ersoy, “A Review and Taxonomy of Activity Recognition on Mobile Phones,” *BioNanoScience*, vol. 3, no. 2, pp. 145–171, jun 2013. [Online]. Available: <http://link.springer.com/10.1007/s12668-013-0088-3>
- [48] A. Bayat, M. Pomplun, and D. A. Tran, “A Study on Human Activity Recognition Using Accelerometer Data from Smartphones,” *Procedia Computer Science*, vol. 34, no. C, pp. 450–457, 2014. [Online]. Available: <http://dx.doi.org/10.1016/j.procs.2014.07.009http://linkinghub.elsevier.com/retrieve/pii/S1877050914008643>
- [49] J. R. Kwapisz, G. M. Weiss, and S. A. Moore, “Cell phone-based biometric identification,” in *2010 Fourth IEEE International Conference on Biometrics: Theory, Applications and Systems (BTAS)*. IEEE, sep 2010, pp. 1–7. [Online]. Available: <http://ieeexplore.ieee.org/document/5634532/>
- [50] E. Miluzzo, A. Varshavsky, S. Balakrishnan, and R. R. Choudhury, “Tap-prints: Your Finger Taps Have Fingerprints,” in *Proceedings of the 10th international conference on Mobile systems, applications, and services - MobiSys '12*. New York, New York, USA: ACM Press, 2012, p. 323. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2307666%}5Cnhttp://dl.acm.org/citation.cfm?doid=2307636.2307666http://dl.acm.org/citation.cfm?doid=2307636.2307666>
- [51] P. Marquardt, A. Verma, H. Carter, and P. Traynor, “(sp)iPhone: Decoding Vibrations From Nearby Keyboards Using Mobile Phone Accelerometers,” in *Proceedings of the 18th ACM conference on Computer and communications security - CCS '11*. New York, New York, USA: ACM Press, 2011, p. 551. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2046771http://dl.acm.org/citation.cfm?doid=2046707.2046771>
- [52] L. Zhang, P. H. Pathak, M. Wu, Y. Zhao, and P. Mohapatra, “AccelWord: Energy Efficient Hotword Detection through Accelerometer,” in *Proceedings of the 13th Annual International Conference on Mobile Systems, Applications, and Services - MobiSys '15*. New York, New York, USA: ACM Press, 2015, pp. 301–315. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=2742647.2742658>
- [53] Y. Michalevsky, D. Boneh, and G. Nakibly, “Gyrophone: Recognizing Speech from Gyroscope Signals,” *Usenix Security*, pp. 1053–1067, 2014. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity14/technical-sessions/presentation/michalevsky>
- [54] Seong-Eun Kim, Yong Kim, Jihyun Yoon, and Eung Sun Kim, “Indoor positioning system using geomagnetic anomalies for smartphones,” in *2012 International Conference on Indoor*

- Positioning and Indoor Navigation (IPIN)*, no. November. IEEE, nov 2012, pp. 1–5. [Online]. Available: <http://ieeexplore.ieee.org/document/6418947/>
- [55] E. Le Grand and S. Thrun, “3-Axis magnetic field mapping and fusion for indoor localization,” in *2012 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI)*, no. 2. IEEE, sep 2012, pp. 358–364. [Online]. Available: <http://ieeexplore.ieee.org/document/6343024/>
- [56] K. P. Subbu, B. Gozick, and R. Dantu, “LocateMe: Magnetic-Fields-Based Indoor Localization Using Smartphones,” *ACM Transactions on Intelligent Systems and Technology*, vol. 4, no. 4, pp. 1–27, sep 2013. [Online]. Available: <http://doi.acm.org/10.1145/2508037.2508054><http://dl.acm.org/citation.cfm?doid=2508037.2508054>
- [57] H. Xia, X. Wang, Y. Qiao, J. Jian, and Y. Chang, “Using Multiple Barometers to Detect the Floor Location of Smart Phones with Built-in Barometric Sensors for Indoor Positioning,” *Sensors*, vol. 15, no. 4, pp. 7857–7877, mar 2015. [Online]. Available: <http://www.mdpi.com/1424-8220/15/4/7857>
- [58] H. Ye, T. Gu, X. Tao, and J. Lu, “Scalable floor localization using barometer on smartphone,” *Wireless Communications and Mobile Computing*, vol. 16, no. 16, pp. 2557–2571, nov 2016. [Online]. Available: <http://eprints.soton.ac.uk/266684/><http://doi.wiley.com/10.1002/wcm.2706>
- [59] W. Falcon and H. Schulzrinne, “Predicting Floor-Level for 911 Calls with Neural Networks and Smartphone Sensor Data,” no. June 2017, pp. 1–16, oct 2017. [Online]. Available: <http://arxiv.org/abs/1710.11122>
- [60] P. Zhou, Y. Zheng, Z. Li, M. Li, and G. Shen, “IODetector: A Generic Service for Indoor Outdoor Detection,” in *Proceedings of the 10th ACM Conference on Embedded Network Sensor Systems - SenSys '12*. New York, New York, USA: ACM Press, 2012, p. 113. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=2426656.2426668>
- [61] H. G. Kayacik, M. Just, L. Baillie, D. Aspinall, and N. Micallef, “Data Driven Authentication: On the Effectiveness of User Behaviour Modelling with Mobile Device Sensors,” *Proceedings of the 3rd Workshop on Mobile Security Technologies (MoST) 2014*, oct 2014. [Online]. Available: <http://arxiv.org/abs/1410.7743>
- [62] N. Micallef, H. G. Kayacik, M. Just, L. Baillie, and D. Aspinall, “Sensor use and usefulness: Trade-offs for data-driven authentication on mobile devices,” in *2015 IEEE International Conference on Pervasive Computing and Communications (PerCom)*. IEEE, mar 2015, pp. 189–197. [Online]. Available: <http://ieeexplore.ieee.org/document/7146528/>
- [63] 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*. New York, New York, USA: ACM Press, 2014, pp. 51–62. [Online]. Available: <http://arxiv.org/abs/1405.3760><http://dl.acm.org/citation.cfm?doid=2666620.2666622>
- [64] L. Li, P. Hu, C. Peng, G. Shen, and F. Zhao, “Epsilon: A Visible Light Based Positioning System,” *USENIX Symposium on Network Systems Design and Implementation*, no. 1, pp. 1–13, 2014. [Online]. Available: <http://panhu.me/pdf/Epsilon.pdf>

- [65] S. E. Z. Mazilu and G. E. Z. Tröster, “A Study on Using Ambient Sensors from Smartphones for Indoor Location Detection,” *12th Workshop On positioning, navigation and communication (WPNC)*. IEEE., 2015.
- [66] D. R. J. Bassett, B. E. Ainsworth, S. R. Leggett, C. A. Mathien, J. A. Main, D. C. Hunter, and G. E. Duncan, “Accuracy of five electronic pedometers for measuring distance walked,” *Medicine & Science in Sports & Exercise*, vol. 28, no. 8, pp. 1071–1077, 1996.
- [67] C. Tudor-Locke, J. E. Williams, J. P. Reis, and D. Pluto, “Utility of Pedometers for Assessing Physical Activity,” *Sports Medicine*, vol. 32, no. 12, pp. 795–808, 2002. [Online]. Available: <http://link.springer.com/10.2165/00007256-200232120-00004>
- [68] S. Crouter, P. L. Schneider, M. Karabulut, and D. R. J. Bassett, “Validity of 10 Electronic Pedometers for Measuring Steps, Distance, and Energy Cost,” *Medicine & Science in Sports & Exercise*, vol. 35, no. 8, pp. 1455–1460, aug 2003. [Online]. Available: <https://insights.ovid.com/crossref?an=00005768-200308000-00030>
- [69] J. Chon and Hojung Cha, “LifeMap: A Smartphone-Based Context Provider for Location-Based Services,” *IEEE Pervasive Computing*, vol. 10, no. 2, pp. 58–67, apr 2011. [Online]. Available: <http://ieeexplore.ieee.org/document/5686873/>
- [70] J. H. Kang, W. Welbourne, B. Stewart, and G. Borriello, “Extracting places from traces of locations,” *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 9, no. 3, p. 58, jul 2005. [Online]. Available: <http://portal.acm.org/citation.cfm?doid=1094549.1094558>
- [71] D. Christin, A. Reinhardt, S. S. Kanhere, and M. Hollick, “A survey on privacy in mobile participatory sensing applications,” *Journal of Systems and Software*, vol. 84, no. 11, pp. 1928–1946, nov 2011. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S0164121211001701>
- [72] C. Welch, “Google took down over 700,000 bad Android apps in 2017,” 2018. [Online]. Available: <https://www.theverge.com/2018/1/30/16951996/google-android-apps-removed-security-2017> (Accessed: June 14, 2018).
- [73] A. Sulleyman, “Millions of people installed infected apps that criminals could use to wreak havoc’ on smartphones,” 2018. [Online]. Available: <https://www.independent.co.uk/life-style/gadgets-and-tech/news/android-apps-google-play-five-nights-survival-guide-adultswine-malware-smartphones-checkpoint-a8159916.html> (Accessed: June 14, 2018).
- [74] Androidrank, “Androidrank,” 2018. [Online]. Available: <https://www.androidrank.org/> (Accessed: June 14, 2018).
- [75] Google, “Google Play Store,” 2018. [Online]. Available: <https://play.google.com> (Accessed: June 14, 2018).
- [76] EDPB, “GDPR Certification Mechanisms, Seals, and Marks,” 2018. [Online]. Available: https://edpb.europa.eu/our-work-tools/accountability-tools/certification-mechanisms-seals-and-marks}_{en (Accessed: October 19, 2018).