

Simon N. Foley (Ed.)

LNCS 11559

Data and Applications Security and Privacy XXXIII

33rd Annual IFIP WG 11.3 Conference, DBSec 2019
Charleston, SC, USA, July 15–17, 2019
Proceedings



ifip



Springer

Simon N. Foley (Ed.)

Data and Applications Security and Privacy XXXIII

33rd Annual IFIP WG 11.3 Conference, DBSec 2019
Charleston, SC, USA, July 15–17, 2019
Proceedings

Editor
Simon N. Foley
Norwegian University of Science
and Technology
Gjøvik, Norway

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-030-22478-3 ISBN 978-3-030-22479-0 (eBook)
<https://doi.org/10.1007/978-3-030-22479-0>

LNCS Sublibrary: SL3 – Information Systems and Applications, incl. Internet/Web, and HCI

© IFIP International Federation for Information Processing 2019

The chapter “Is My Phone Listening in? On the Feasibility and Detectability of Mobile Eavesdropping” is Open Access. This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>). For further details see license information in the chapter.

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, expressed or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland



Is My Phone Listening in? On the Feasibility and Detectability of Mobile Eavesdropping

Jacob Leon Kröger^{1,2(✉)} and Philip Raschke¹

¹ Technische Universität Berlin, Berlin, Germany
{kroeger, philip.raschke}@tu-berlin.de

² Weizenbaum Institute for the Networked Society, Berlin, Germany

Abstract. Besides various other privacy concerns with mobile devices, many people suspect their smartphones to be secretly eavesdropping on them. In particular, a large number of reports has emerged in recent years claiming that private conversations conducted in the presence of smartphones seemingly resulted in targeted online advertisements. These rumors have not only attracted media attention, but also the attention of regulatory authorities. With regard to explaining the phenomenon, opinions are divided both in public debate and in research. While one side dismisses the eavesdropping suspicions as unrealistic or even paranoid, many others are fully convinced of the allegations or at least consider them plausible. To help structure the ongoing controversy and dispel misconceptions that may have arisen, this paper provides a holistic overview of the issue, reviewing and analyzing existing arguments and explanatory approaches from both sides. Based on previous research and our own analysis, we challenge the widespread assumption that the spying fears have already been disproved. While confirming a lack of empirical evidence, we cannot rule out the possibility of sophisticated large-scale eavesdropping attacks being successful and remaining undetected. Taking into account existing access control mechanisms, detection methods, and other technical aspects, we point out remaining vulnerabilities and research gaps.

Keywords: Privacy · Smartphone · Eavesdropping · Spying · Listening · Microphone · Conversation · Advertisement

1 Introduction

Smartphones are powerful tools that make our lives easier in many ways. Since they are equipped with a variety of sensors, store large amounts of personal data and are carried throughout the day by many people, including in highly intimate places and situations, they also raise various privacy concerns.

One widespread fear is that smartphones could be turned into remote bugging devices. For years, countless reports have been circulating on the Internet from people who claim that things they talked about within earshot of their phone later appeared in targeted online advertisements, leading many to believe that their private conversations must have been secretly recorded and analyzed.

The reported suspicious ads range across many product and service categories, including clothing, consumer electronics, foods and beverages, cars, medicines, holiday destinations, sports equipment, pet care products, cosmetics, and home appliances – and while some of these ads were described as matching an overall discussion topic, others allegedly promoted a brand or even a very specific product mentioned in a preceding face-to-face conversation [6, 12]. Some people claim to have experienced the phenomenon frequently and that they have successfully reproduced it in private experiments. Interestingly, many of the purported witnesses emphasize that the advertised product or service seems not related to places they have visited, terms they have searched for online, or things they have mentioned in text messages, emails or social media [6, 40]. Furthermore, some reports explicitly rate it as unlikely that the respective advertisements were selected by conventional targeting algorithms, as they lay notably outside the range of advertising normally received and did sometimes not even appear to match the person's consumer profile (e.g. in terms of interests, activities, age, gender, or relationship status) [6, 41].

Numerous popular media outlets have reported on these alleged eavesdropping attacks [3]. In a *Forbes* article, for instance, the US-based market research company Forrester reports that at least 20 employees in its own workforce have experienced the phenomenon for themselves [40]. The same holds true for one in five Australians, according to a recent survey [38]. Even the US House Committee on Energy and Commerce has started to investigate the issue by sending letters to Google and Apple inquiring about the ways in which iOS and Android devices record private conversations [77].

Many commentators, including tech bloggers, researchers and business leaders, on the other hand, view the fear that private companies could target their ads based on eavesdropped conversations as baseless and paranoid. The reputational risk, it is argued, would be far too high to make this a viable option [76]. With regard to CPU, battery and data storage limitations, former Facebook product manager Antonio García Martínez even considers the alleged eavesdropping scenario to be economically and technically unfeasible [51]. As an alternative explanation for suspiciously relevant ads, he points to the many established and well-documented methods that companies successfully use to track, profile and micro-target potential customers. Yet another possible explanation states that the frequently reported phenomenon is merely a product of chance, potentially paired with some form of confirmation bias [41]. Finally, some commentators also suggest that topics of private conversations are sometimes inspired by unconsciously processed advertisements, which may later cause the perception of being spied upon when the respective ad is encountered again [28].

Many views, theories and arguments have been put forward in attempt to explain the curious phenomenon, including experimental results and positions from the research community. However, a consensus has not yet been reached, not even regarding the fundamental technical feasibility of the alleged eavesdropping attacks. Therefore, this paper reviews, verifies and compares existing arguments from both sides of the discourse. Apart from providing a structured overview of the matter, conclusions about the feasibility and detectability of smartphone-based eavesdropping are drawn based on existing research and our own analysis.

In accordance with the reports found on the phenomenon, this paper will focus on smartphones – specifically, iOS and Android devices. Since smartphones are the most widespread consumer electronics device, and since iOS and Android together clearly dominate the mobile OS market [70], this choice seems justified to us. However, most of the considerations in this paper are applicable to other types of mobile devices and other operating systems as well.

The remainder of this paper is structured as follows. In Sect. 2, we describe the underlying threat model, distinguishing between three possible adversaries. Section 3 examines the possibility of using smartphone microphones for stealthy eavesdropping, expanding on aspects of security permissions and user notifications. Similarly, Sect. 4 considers smartphone motion sensors as a potential eavesdropping channel, taking into account sampling frequency limits enforced by mobile operating systems. Section 5 then looks into the effectiveness of existing mitigation and detection techniques developed by Google, Apple, and the global research community. In Sect. 6, the ecosystem providers themselves are considered as potential adversaries. Section 7 evaluates the technical and economic feasibility of large-scale eavesdropping attacks. After that, Sect. 8 examines ways in which governmental and criminal hackers can compromise the speech privacy of smartphone users. Finally, Sect. 9 provides a discussion of analysis results, followed by a conclusion in Sect. 10.

2 Threat Model

To target advertisements based on smartphone eavesdropping, an organization A, who is responsible for selecting the audience for certain online ads (either the advertiser itself or a contractor entrusted with this task, such as an advertising network¹), needs to somehow gain access to sensor data² from the corresponding mobile device, or to information derived from the sensor data.

Initially, speech is recorded through the smartphone by an actor B, which could be either (1) the operating system provider itself, e.g. Apple or Google, (2) non-system apps installed on the device, or (3) third-party libraries³ included in these apps. Potentially after some processing and filtering, which can happen locally on the device or on remote servers, actor B shares relevant information extracted from the recording – directly or through intermediaries – with organization A (unless A and B are one and the same actor, which is also possible).

Organization A then uses the received information to identify the smartphone owner as a suitable target for specific ads and sends a corresponding broadcast request to an ad publisher (organization A could also publish the ads itself if it has access to ad distribution channels). Finally, the publisher displays the ads on websites or apps – either on the smartphone through which the speech was recorded or on other devices

¹ Advertising networks are companies that match demand and supply of online ad space by connecting advertisers to ad publishers. They often hold extensive amounts of data on individual internet users to enable targeted advertising [17].

² “sensor data” can refer to either audio recordings or motion sensor data (see Sects. 3, 4).

³ The role and significance of third-party apps will be further explained in Sect. 3.1.

that can be linked⁴ to the smartphone owner, for example through logins, browsing behavior, or IP address matching. The websites and apps on which the advertisements appear do not reveal who recorded the smartphone owner's speech. Not even organization A necessarily understands how and by whom the received profiling information was initially collected. For illustration, Fig. 1 presents a simplified overview of the threat model.

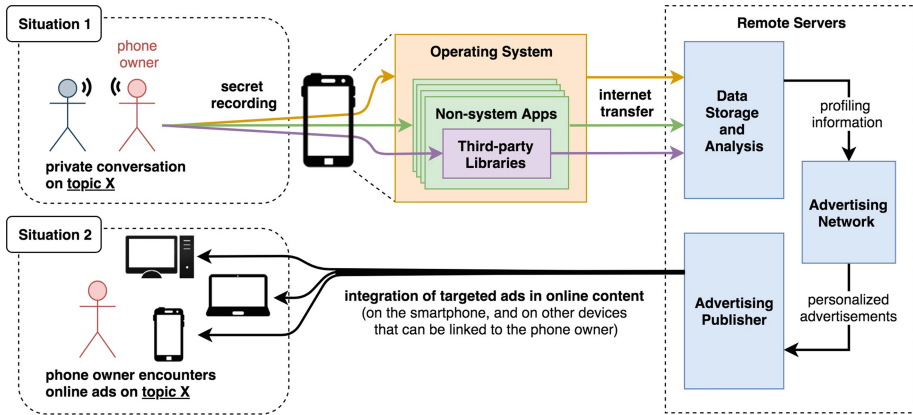


Fig. 1. A schematic and simplified overview of the threat model.

3 Microphone-Based Eavesdropping

Modern smartphones have the capability to tape any sort of ambient sound through built-in microphones, including private conversations, and to transmit sensitive data, such as the recording itself or information extracted from recorded speech, to remote servers over the Internet. Mobile apps installed on a phone could exploit these capabilities for secret eavesdropping. Aspects concerning app permissions and user notifications that could affect the feasibility and visibility of such an attack are examined in the following two subsections.

3.1 Microphone Access Permission

Before an app can access microphones in Android and iOS devices, permission has to be granted by the user. However, people tend to accept such requests blindly if they are interested in an app's functionality [10]. A survey of 308 Android users found that only 17% of respondents paid attention to permissions during app installation, and no more than 3% of the participants correctly answered the related comprehension questions [24].

⁴ For more information on cross-device tracking, refer to [65].

Encouraging app development at the expense of user privacy, current permission systems are much less strict than they were in early smartphones and have been criticized as “coarse grained and incomplete” [59]. Also, once a permission is granted, it is usually not transparent for users when and for which particular purpose data is being collected and to which servers it is being sent [62].

To include analytics and advertising capabilities, apps commonly make use of third-party libraries, i.e., code written by other companies. These libraries share multimedia permissions, such as microphone access, with their corresponding host app and are often granted direct Internet access [39]. Apart from the concern that third-party libraries are easily over-privileged, it is considered problematic that app developers often have limited or no understanding of the library code, which can also be changed dynamically at runtime [59]. Thus, not only users but also app developers themselves may be unaware of privacy leaks based on the abuse of granted permissions.

A large variety of existing apps has access to smartphone microphones. Examining over 17.000 popular Android apps, Pan et al. found that 43.8% ask for permission to record audio [59].

3.2 User Notifications and Visibility

Android and iOS apps with microphone permission can not only record audio at any time while they are active, i.e. running in the foreground, but also while they are in background mode, under certain conditions [7, 31]. Background apps have limited privileges and are often suspended to conserve the device’s limited resources. In cases, however, where they request the system to stay alive and continue recording while not in the foreground, there are ways to indicate this to the user.

In iOS, the status bar will automatically turn bright red when recording takes place in the background, allowing the user to immediately detect potentially unwanted microphone activity [8].

While the latest release of Android (version 9 Pie) implements similar measures [31], some older versions produce no visible indication when background apps access the microphone [10]. In this context, it might be worth noting that Android has been widely criticized for its slow update cycle, with hundreds of millions of devices running on massively outdated versions [56]. Also, quite obviously, notifications in the graphical user interface are only visible as long as the device’s screen is not turned off. And finally, some experimenters have already succeeded in circumventing the notification requirements for smartphone media recordings [69].

4 Motion Sensor-Based Eavesdropping

Adversaries might be able to eavesdrop on conversations through cell phones without accessing the microphone. Studies have shown that smartphone motion sensors – more specifically, accelerometers and gyroscopes – can be sensitive enough to pick up sound vibrations and possibly even reconstruct speech signals [36, 54, 79].

4.1 Experimental Research Findings

There are opposing views on whether non-acoustic smartphone sensors capture sounds at normal conversational loudness. While Anand and Saxena did not notice an apparent effect of live human speech on motion sensors in several test devices [3], other studies report very small but measurable effects of machine-rendered speech, significant enough to reconstruct spoken words or phrases [54, 79].

Using only smartphone gyroscopes, researchers from Israel's defense technology group Rafael and Stanford University were able to capture acoustic signals rich enough to identify a speaker's gender, distinguish between different speakers and, to some extent, track what was being said [54]. In a similar experiment, Zhang et al. demonstrated the feasibility of inferring spoken words from smartphone accelerometer readings in real-time, even in the presence of ambient noise and user mobility [79]. According to their evaluation, the achieved accuracies were comparable to microphone-based hotword detection applications such as Samsung S Voice and Google Now.

Both [79] and [54] have notable limitations. First of all, their algorithms were only able to detect a small set of predefined keywords instead of performing full speech recognition. Also, the speech in both experiments was produced by loudspeakers or phone speakers, which may result in acoustic properties different from live human speech. In [54], the playback device and the recording smartphone even shared a common surface, leading critics to suggest that the observed effect on sensor readings was not caused by aerial sound waves, but rather by direct surface vibrations [3]. Also, in contrast to Zhang et al., this approach only achieved low recognition accuracies, particularly for speaker-independent hotword detection. By their own admission, however, the authors of [54] are "security experts, not speech recognition experts" [32]. Therefore, the study should be regarded as an initial exploration rather than a perfect simulation of state-of-the-art spying techniques. With regard to the effectiveness of their approach, the researchers pointed out several possible directions for future improvement.

It might also be noteworthy that patents have already been filed for methods to capture acoustic signals through motion sensors, including a "method of detecting a user's voice activity using an accelerometer" [21] and a "system that uses an accelerometer in a mobile device to detect hotwords" [55].

4.2 Sampling Frequency Limits

In order to limit energy consumption and because typical applications of smartphone motion sensors do not require highly sampled data, current mobile operating systems impose a cap on the sampling frequency of motion sensors, such as a maximum of 200 Hz for accelerometer readings in Android [3] and 100 Hz for gyroscopes in iOS [32]. For comparison, the fundamental frequency of the human speaking voice typically lies between 85 Hz and 155 Hz for men and 165 Hz and 255 Hz for women [79]. Thus, if at all, non-acoustic smartphone sensors can only capture a limited range of speech sounds, which presents a challenge to speech reconstruction attacks.

With the help of the aliasing effect explained in [54], however, it is possible to indirectly capture tones above the enforced frequency limits. Furthermore, experiments show that motion sensor signals from multiple co-located devices can be merged to obtain a signal with increased sampling frequency, significantly improving the effectiveness of speech reconstruction attacks [36]. Two or more smartphones that are located in proximity to each other and whose sensor readings are shared – directly or indirectly – with the same actor may therefore pose an increased threat to speech privacy.

It should also be noted that motion sensors in smartphones are usually capable of delivering much higher sampling frequencies (often up to 8 kHz) than the upper bounds prescribed by mobile operating systems [3]. Researchers already expressed concern that adversaries might be able to override and thereby exceed the software-based limits through patching applications or kernel drivers in mobile devices [3, 54].

4.3 Sensor Access Permissions and Energy Efficiency

While certain hardware components, such as camera, microphone and the GPS chip, are typically protected by permission mechanisms in mobile operating systems, motion sensors can be directly accessed by third-party apps in iOS and Android without any prior notification or request to the user [32, 45]. Thus, there is usually no way for smartphone owners to monitor, let alone control when and for what purposes data from built-in accelerometers and gyroscopes is collected. Even visited websites can often access smartphone motion sensors [32]. Exploiting accelerometers and gyroscopes to intrude user privacy is also much more energy-efficient and thus less conspicuous than recording via microphone [79].

5 Existing Mitigation and Detection Techniques

Many methods are applied by ecosystem providers and security researchers to screen mobile apps for vulnerabilities and malicious behavior. The following two subsections examine existing efforts with regard to their potential impact on the feasibility and detectability of mobile eavesdropping attacks.

5.1 App Inspections Conducted by Ecosystem Providers

Both iOS and Android apply a combination of static, dynamic and manual analysis to scan new and existing apps on their respective app market for potential security threats and to ensure that they operate as advertised [78]. Clearly, as the misbehavior of third-party apps can ultimately damage their own reputation, the platforms have strong incentives to detect and prevent abuse attempts.

Nevertheless, countless examples of initially undetected malware and privacy leaks have shown that the security screenings provided by Google and Apple are not always successful [19]. Google Play’s app inspection process has even been described as “fundamentally vulnerable” [29]. In a typical cat-and-mouse game, malicious apps evolve quickly to bypass newly implemented security measures [63], sometimes by

using “unbearably simple techniques” [29]. In Android devices from uncertified manufacturers, malware may even be pre-installed before shipment [14]. Significant vulnerabilities have also been found in official built-in apps. Apple’s FaceTime app, for example, allowed potential attackers to gain unauthorized access to iPhone cameras and microphones without any requirement of advanced hacking skills [15].

Leaving security loopholes aside, the existing security mechanisms do not guarantee privacy protection in terms of data minimization and transparency. Many mobile apps collect personal data with no apparent relevance to the advertised functionality [18, 62]. Even well-known apps like Uber have not been prevented from collecting sensitive user data that is not required for the service they offer [46].

There are also many documented cases of mobile apps using their microphone access in unexpected ways. An example that has received a lot of media attention recently is the use of so-called “ultrasonic beacons”, i.e. high-pitched Morse-style audio signals inaudible to the human ear which are secretly played in stores or embedded in TV commercials and other broadcast content in order to be able to unobtrusively track the location, activities and media consumption habits of consumers [10]. For this to work, the data subject needs to carry a receiving device that records and scans ambient sound for relevant ultrasonic signals and sends them back to the tracking network for automated comparison. A constantly growing number of mobile apps – several hundred already, some of them very popular – are using their microphone permission for exactly that purpose, often without properly informing the user about it [10, 47]. These apps, some of which are targeted at children and would not require audio recording for their core functionality, may even detect sounds while the phone is locked and carried in a pocket [47]. Even in cases where users are aware that their phone listens in, it is not clear to them what the audio stream is filtered for exactly and what information is being exfiltrated. Thus, the example of ultrasonic beacons shows how apps that have been approved into Apple’s App Store and Google Play can exploit their permissions for dubious and potentially unexpected tracking purposes.

Finally, it should not be overlooked that smartphone apps can also be obtained from various non-official sources, circumventing Apple’s and Google’s permission systems and auditing processes [62]. In Android, users are free in choosing the source of their applications [78]. Following a more restrictive policy, iOS only allows users to install apps downloaded from the official Apple App Store. However, kernel patches can be used to gain root access and remove software restrictions in iOS (“iOS jailbreaking”), which enables users to install apps from uncertified publishers [62].

5.2 App Inspections Conducted by the Research Community

In addition to the checks conducted by Google and Apple, mobile apps are being reviewed by a broad community of security and privacy researchers. A wide and constantly expanding range of manual and automated methods is applied for this purpose.

Pan et al., for instance, scanned 17,260 popular Android apps from different app markets for potential privacy leaks [59]. Through examining their media permissions, privacy policies and outgoing network flows, the researchers tried to identify apps that upload audio recordings to the Internet without explicitly informing the user about it.

While unveiling other serious forms of privacy violations, they found no evidence of such behavior. Based on these findings, the widely held suspicion of companies secretly eavesdropping on smartphone users was already portrayed as refuted in news headlines [34, 80].

However, the study comes with numerous limitations: Apart from considering only a small fraction of the over 2 million available Android apps, the researchers did not examine media exfiltration from app background activity, did not consider the use of privileged APIs, only tested a limited amount of each app’s functionalities for a short amount of time, used a controlled test environment with no real human interactions, did not consider iOS apps at all, and were not able to detect media that was intentionally obfuscated, encrypted at the application-layer, or sent over the network in non-standard encoding formats. Perhaps most importantly, Pan et al. were not able to rule out the scenario of apps transforming audio recordings into less detectable text transcripts or audio fingerprints before sending the information out. This would be a very realistic attack scenario. In fact, various popular apps are known to compress recorded audio in such a way [10, 33]. While all the choices that Pan et al. made regarding their experimental setup and methodology are completely understandable and were communicated transparently, the limitations do limit the significance of their findings. All in all, their approach would only uncover highly unsophisticated eavesdropping attempts.

Of course, many other researchers have also tried to detect privacy leaks in iOS and Android apps [62]. Besides analyzing decompiled code, permission requests and generated network traffic, other factors, such as battery power consumption and device memory usage, can also be monitored to detect suspicious app behavior [67]. Although some experts claim to have observed certain mobile apps recording and sending out audio with no apparent justification [58], the scientific community has not yet produced any hard evidence for large-scale eavesdropping through smartphone microphones.

Like the above-cited work by Pan et al., however, other existing methods to identify privacy threats in mobile devices also come with considerable limitations. Due to its closed-source nature, there is generally a lack of scalable tools for detecting malicious apps within iOS [19]. While, on the other hand, numerous efficient methods have been proposed for automatically scanning Android apps, none of these approaches is totally effective at detecting privacy leaks [59]. As with security checks of the official app stores (see Sect. 5.1), there is a wide range of possible obfuscation techniques and covert channels to circumvent detection mechanisms developed by the scientific community [10, 67]. Furthermore, many of the existing approaches do not indicate if detected data exfiltration activities are justified with regard to an app’s advertised functionality [62]. Yerukhimovich et al. even suggest that apps classified as safe or non-malicious are more likely to leak private information than typical “malware” [78].

Therefore, the fact that no evidence for large-scale mobile eavesdropping has been found so far should not be interpreted as an all-clear. It could only mean that it is difficult – under current circumstances perhaps even impossible – to detect such attacks effectively.

6 Ecosystem Providers as Potential Adversaries

Not only third-party apps but also mobile operating systems themselves can access privacy-sensitive smartphone data and transfer it over the Internet. It has been known for years that both, iOS and Android, do so extensively [5]. Examining the amount of data sent back to Google's and Apple's servers from test devices, a recent study found that iPhones – on average – received four requests per hour from their manufacturer during idle periods, and eighteen requests during periods of heavy use [68]. Leaving these numbers far behind, Android phones received forty hourly requests from Google when in idle state and ninety requests during heavy use. Of course, the number of requests per hour has only limited informational value. Data is often collected much more frequently, such as on a secondly basis or even constantly, to be later aggregated, compressed and sent out in data bundles [5].

While the establishment of network connections can be monitored, many aspects of data collection and processing in smartphones remain opaque. The source code of iOS is not made publicly available, and while Android is based on code from the Android Open Source Project, several of Google's proprietary apps and system components are closed-source as well [2]. Due to the resulting lack of transparency, it cannot be reliably ruled out that sensitive data is collected and processed without the will or knowledge of the smartphone owner – although, naturally, this would represent a considerable legal and reputational risk for the corresponding platform provider.

As an intermediary between applications and hardware resources, operating systems control the access to smartphone sensors, including microphones, accelerometers and gyroscopes, and can also decide whether or not sensor activity is indicated to the user on the device's screen. Other than with third-party apps, there is no superior authority in the system supervising the actions and decisions of iOS and Android. While external security experts can carry out inspections using similar methods as outlined in Sect. 5.2, they also face similar limitations. There is no reason to assume that operating systems refrain from using sophisticated obfuscation techniques to conceal their data collection practices. Additionally, being in control of the whole system, iOS and Android can access data on different levels of their respective software stack, which gives them more options for stealthy data exfiltration and could possibly impede detection.

7 Technical and Economic Feasibility

Even where adversaries manage to get around security measures and evade detection, it remains questionable whether a continuous and large-scale eavesdropping operation for the purpose of ad targeting would be technically feasible and economically viable. Based on estimations of CPU, battery, network transfer and data storage requirements, some commentators already stated their conclusion that such an operation would be far too expensive [51, 76] and may “strain even the resources of the NSA” [71]. Taking into account their underlying assumptions, these estimates appear valid. However, there are several ways in which smartphone-based eavesdropping could be made much more efficient and scalable, including:

- **Low quality audio recording.** To reduce the required data storage, processing power and energy consumption, adversaries could record audio at low bitrates. Speech signals do not even have to be intelligible to the human ear to be recognized and transcribed into text by algorithms [54].
- **Local pre-processing.** Some steps in the processing of recordings (e.g. transcription, extraction of audio features, data filtering, keyword matching, compression) can be performed locally on the device in order to transmit only the most relevant data to remote servers and thus reduce network traffic and required cloud storage.
- **Keyword detection instead of full speech recognition.** The amounts of processing power required for automatic speech recognition can be prohibitively high for local execution on mobile devices. A less CPU-intensive alternative to full speech recognition is keyword detection, where only a pre-defined vocabulary of spoken words is recognized. Such systems can even run on devices with much lower computational power than smartphones, such as 16-bit microcontrollers [25]. It has been argued that it would still be too taxing for mobile devices to listen out for the “millions or perhaps billions” of targetable keywords that could potentially be dropped in private conversations [51]. However, instead of listening for specific product and brand names, audio recordings can simply be scanned for trigger words that indicate a person’s interest, such as “love”, “enjoyed”, or “great”, in order to identify relevant snippets of the recording, which can then be analyzed in more depth. In fact, this very audio analysis method has already been patented, with the specific declared purpose of informing “targeted advertising and product recommendations” [22].
- **Selective recording.** Instead of recording continuously, an adversary could only record at selected moments using wake words or triggers based on time, location, user activity, sound level, and other context variables. This could significantly reduce the amount of required storage and network traffic [67].

Mobile apps that use all or some of the above techniques can be light enough to run smoothly on smartphones, as numerous commercial apps and research projects show [9, 10, 33, 58, 67].

But even if it is possible for companies to listen in on private conversations, some argue that this information might not be of much value to advertisers, since they would need to know a conversation’s context and speaker personalities very well in order to accurately infer personal preferences and purchase intentions from spoken phrases [51]. This argument is reasonable, but can equally be applied to many other profiling methods, including online tracking and location tracking, which are widely used nonetheless. Of course, where contextual information is sparse, such methods may lead to wrong conclusions about the respective data subject, possibly resulting in poor and inefficient ad targeting. However, this would not conflict with the above-mentioned reports of suspected eavesdropping: While the ads were perceived as inspired by topics raised in private conversations, they did not always reflect the purported witnesses’ actual needs and wants [6, 12].

From an outside perspective, it cannot be precisely determined how profitable certain types of personal data are for advertisers. It is therefore difficult, if not impossible, to draw up a meaningful cost-benefit calculation. However, it can generally

be assumed that private conversations contain a lot of valuable profiling information, especially when speakers express their interest in certain products or services. It is also worth mentioning that some of the world's largest companies earn a significant portion of their revenue through advertising – for Google and Facebook, this portion amounted to 85% and 98% in 2018, respectively [1, 23]. Profits from advertising can be considerably increased through effective targeting, which requires the collection of detailed personal information [68]. There is no doubt that smartphone sensor data can be very useful for this purpose. A recently filed patent describes, for example, how “local signals” from a mobile device, including motion sensor data and audio data from the microphone, can be analyzed to personalize a user's Facebook news feed [50].

8 Unauthorized Access to Smartphones

Although this is most likely no explanation for suspicious ad placement, it should be noted that there are many ways in which skilled computer experts or “hackers” can gain unauthorized access to mobile devices. The widespread use of smartphones makes them a particularly attractive hacking target [4].

Not only cyber criminals, but also law enforcement agencies and secret services invest heavily in their capabilities to exploit software flaws and other security vulnerabilities in consumer electronics [73]. It has been known for some time that intelligence agencies, such as NSA, GCHQ, and CIA, are equipped with tools to secretly compromise devices running iOS, Android and other mobile operating systems, enabling them “to move inside a system freely as if they owned it” [66, 75].

In addition to accessing sensitive data, such as geo-location, passwords, personal notes, contacts, and text messages, this includes the ability to turn on a phone's microphone without a user's consent or awareness [11]. With the help of specialized tools, smartphone microphones can even be tapped when the device is (or seems) switched off [73]. Such attacks can also be successful in high-security environments. In a recent case, for example, more than 100 Israeli servicemen had their phones infected with spyware that allowed unknown adversaries to control built-in cameras and microphones [57].

Besides the United States and some European nations, other developed countries, such as Russia, Israel and China, also have highly sophisticated spying technology at their disposal [75]. Less developed countries and other actors can buy digital eavesdropping tools from a flourishing industry of surveillance contractors at comparatively low prices [60]. That not only secret services but also law enforcement agencies in the US can be authorized to convert smartphones into “roving bugs” to listen in on private conversations has been confirmed in a 2012 court ruling [16]. Eavesdropping capabilities of criminal organizations should not be underestimated, either. According to a report by McAfee and the Center for Strategic and International Studies (CSIS), there are 20 to 30 cybercrime groups with “nation-state level” capacity in countries of the former Soviet Union alone [52].

9 Discussion

So far, despite significant research efforts, no evidence has been found to confirm the widespread suspicion that firms are secretly eavesdropping on smartphone users to inform ads. To the best of our knowledge, however, the opposite has not been proven either. While some threat scenarios (e.g. the constant transfer of uncompressed audio recordings into the cloud) can be ruled out based on existing security measures and considerations regarding an attack's visibility, cost and technical feasibility, there are still many security vulnerabilities and a fundamental lack of transparency that potentially leave room for more sophisticated attacks to be successful and remain undetected.

In comparison with the researchers cited in this paper, it can be assumed that certain companies have significantly more financial resources, more training data, and more technical expertise in areas such as signal processing, data compression, covert channels, and automatic speech recognition. This is – besides unresolved contradictions between cited studies and large remaining research gaps – another reason why existing work should not be seen as final and conclusive, but rather as an initial exploration of the issue.

While this paper focuses on smartphones, it should be noted that microphones and motion sensors are also present in a variety of other Internet-connected devices, including not only VR headsets, wearable fitness trackers and smartwatches, but also baby monitors, toys, remote controls, cars, household appliances, laptops, and smart speakers. Some of these devices may have weaker privacy safeguards than smartphones. For instance, they may not ask for user permission before turning on the microphone or may not impose a limit on sensor sampling frequencies. Numerous devices, including smart TVs [13], smart speakers [27], and connected toys [26], have already been suspected to spy on private conversations of their users. Certain smart home devices, such as home security alarms, may even contain a hidden microphone without disclosing it in the product specifications [44]. For these reasons, it is essential to also thoroughly examine non-smartphone devices when investigating suspicions of eavesdropping.

It is quite possible, at the same time, that the fears of advertising companies eavesdropping on private conversations are unfounded. Besides the widespread attribution to chance, one alternative approach to explaining strangely accurate advertisements points to all the established tracking technologies commonly employed by advertisers that do not depend on any phone sensors or microphones [51].

Drawing from credit card networks, healthcare providers, insurers, employers, public records, websites, mobile apps, and many other sources, certain multi-national corporations already hold billions of individual data points on consumers' location histories, browsing behaviors, religious and political affiliations, occupations, socioeconomic backgrounds, health conditions, personality traits, product preferences, and so on [17, 64]. Although their own search engines, social networks, email services, route planners, instant messengers, and media platforms already give them intimate insight into the lives of billions of people, advertising giants like Facebook and Google also intensively track user behavior on foreign websites and apps. Of the 17,260 apps examined in [59], for example, 48.22% share user data with Facebook in the

background. Through their analytics services and like buttons, Google and Facebook can track clicks and scrolls of Internet users on a vast number of websites [17].

The deep and potentially unexpected insights that result from such ubiquitous surveillance can be used for micro-targeted advertising and might thereby create an illusion of being eavesdropped upon, especially if the data subject is ill-informed about the pervasiveness and impressive possibilities of data linkage.

Even without being used for audio snooping, smartphones (in their current configuration) allow a large variety of actors to track private citizen in a much more efficient and detailed way than would ever have been possible in even the most repressive regimes and police states of the 20th century. At the bottom line, whether sensitive information is extracted from private conversations or collected from other sources does not make much difference to the possibilities of data exploitation and the entailing consequences for the data subject. Therefore, whether justified or not, the suspicions examined in this paper eventually lead to a very fundamental question: What degree of surveillance should be considered acceptable for commercial purposes like targeted advertising? Although this paper cannot offer an answer to this political question, it should not be forgotten that constant surveillance is by no means a technical necessity and that, by definition, democracies should design and regulate technology to primarily reflect the values of the public, not commercial interests.

Certainly, the fear of eavesdropping smartphones should never be portrayed as completely unfounded, as various criminal and governmental actors can gain unauthorized access to consumer electronics. Although such attacks are unlikely to result in targeted advertisement, they equally deprive the user of control over his or her privacy and might lead to other unpredictable harms and consequences. For example, digital spying tools have been used to infiltrate the smartphones of journalists [49] and human rights activists [60] for repressive purposes.

Finally, it should be recognized that – apart from the linguistic contents of speech – microphones and motion sensors may unexpectedly transmit a wealth of other sensitive information. Through the lens of advanced analytics, a voice recording can reveal a speaker's identity [53], physical and mental health state [20, 37], and personality traits [61], for example. Accelerometer data from mobile devices may implicitly contain information about a user's location [35], daily activities [48], eating, drinking and smoking habits [72, 74], degree of intoxication [30], gender, age, body features and emotional state [43] and can also be used to re-construct sequences of text entered into a device, including passwords [42].

10 Conclusion

After online advertisements seemingly adapted to topics raised in private face-to-face conversations, many people suspect companies to secretly listen in through their smartphones. This paper reviewed and analyzed existing approaches to explaining the phenomenon and examined the general feasibility and detectability of mobile eavesdropping attacks. While it is possible, on the one hand, that the strangely accurate ads were just a product of chance or conventional profiling methods, the spying fears were

not disproved so far, neither by device manufacturers and ecosystem providers nor by the research community.

In our threat model, we considered non-system mobile apps, third-party libraries, and ecosystem providers themselves as potential adversaries. Smartphone microphones and motion sensors were investigated as possible eavesdropping channels. Taking into account permission requirements, user notifications, sensor sampling frequencies, limited device resources, and existing security checks, we conclude that – under the current levels of data collection transparency in iOS and Android – sophisticated eavesdropping operations could potentially be run by either of the above-mentioned adversaries without being detected. At this time, no estimate can be made as to the probability and economic viability of such attacks.

References

1. Alphabet Inc.: Alphabet Announces Fourth Quarter and Fiscal Year 2018 Results (2019). https://abc.xyz/investor/static/pdf/2018Q4_alphabet_earnings_release.pdf?cache=adc3b38
2. Amadeo, R.: Google's iron grip on Android: Controlling open source by any means necessary (2018). <https://arstechnica.com/gadgets/2018/07/googles-iron-grip-on-android-controlling-open-source-by-any-means-necessary/>
3. Anand, S.A., Saxena, N.: Speechless: analyzing the threat to speech privacy from smartphone motion sensors. In: 2018 IEEE Symposium on Security and Privacy, San Francisco, CA, pp. 1000–1017. IEEE (2018). <https://doi.org/10.1109/SP.2018.00004>
4. Aneja, L., Babbar, S.: Research trends in malware detection on Android devices. In: Panda, B., Sharma, S., Roy, N. (eds.) Data Science and Analytics. Communications in Computer and Information Science, vol. 799, pp. 629–642. Springer, Singapore (2018). https://doi.org/10.1007/978-981-10-8527-7_53
5. Angwin, J., Valentino-DeVries, J.: Apple, Google Collect User Data (2011). <https://www.wsj.com/articles/SB10001424052748703983704576277101723453610>
6. Anonymous: YouTube user demonstrating how Facebook listens to conversations to serve ads (2017). https://www.reddit.com/r/videos/comments/79i4cj/youtube_user_demonstrating_how_facebook_listens/
7. Apple: Background Execution. <https://developer.apple.com/library/archive/documentation/iPhone/Conceptual/iPhoneOSProgrammingGuide/BackgroundExecution/BackgroundExecution.html>
8. Apple: Record - iPhone User Guide. <https://help.apple.com/iphone/11/?lang=en#/iph4d2a39a3b>
9. Arcas, B.A., et al.: Now playing: continuous low-power music recognition. arXiv Comput. Res. Repos. abs/1711.10958 (2017). <http://arxiv.org/abs/1711.10958>
10. Arp, D., et al.: Privacy threats through ultrasonic side channels on mobile devices. In: 2017 IEEE European Symposium on Security and Privacy (EuroS&P), Paris, France, pp. 35–47. IEEE (2017). <https://doi.org/10.1109/EuroSP.2017.33>
11. Ball, J.: Angry Birds and “leaky” phone apps targeted by NSA and GCHQ for user data (2014). <https://www.theguardian.com/world/2014/jan/27/nsa-gchq-smartphone-app-angry-birds-personal-data>
12. BBC News Services: Is your phone listening in? Your stories (2017). <https://www.bbc.com/news/technology-41802282>

13. Beres, D.: How To Stop Your Smart TV From Eavesdropping On You (2015). https://www.huffpost.com/entry/your-samsung-tv-is-spying-on-you_n_6647762
14. Bocek, V., Chrysaidos, N.: Android devices ship with pre-installed malware (2018). <https://blog.avast.com/android-devices-ship-with-pre-installed-malware>
15. Bogost, I.: FaceTime Is Eroding Trust in Tech (2019). <https://www.theatlantic.com/technology/archive/2019/01/apple-facetime-bug-you-cant-escape/581554/>
16. Brown, A.J.: United States v. Oliva (United States Court of Appeals, D.C. No. 3:07-cr-00050-BR-1) (2012)
17. Christl, W.: Corporate Surveillance in Everyday Life. Cracked Labs, Vienna (2017)
18. Christl, W., Spiekermann, S.: Networks of Control: A Report on Corporate Surveillance, Digital Tracking, Big Data & Privacy. Facultas, Vienna (2016)
19. Cimitile, A., et al.: Machine learning meets iOS malware: identifying malicious applications on Apple environment. In: Proceedings of the 3rd International Conference on Information Systems Security and Privacy, Porto, Portugal, pp. 487–492. SciTePress (2017). <https://doi.org/10.5220/0006217304870492>
20. Cummins, N., et al.: Speech analysis for health: current state-of-the-art and the increasing impact of deep learning. Methods (2018). <https://doi.org/10.1016/j.ymeth.2018.07.007>
21. Dusan, S.V., et al.: System and Method of Detecting a User's Voice Activity Using an Accelerometer (Patent No.: US9438985B2) (2014). <https://patents.google.com/patent/US9438985B2/en>
22. Edara, K.K.: Keyword Determinations from Voice Data (Patent No.: US20140337131A1) (2014). <https://patents.google.com/patent/US20140337131A1/en>
23. Facebook: Facebook Reports Fourth Quarter and Full Year 2018 Results. https://s21.q4cdn.com/399680738/files/doc_financials/2018/Q4/Q4-2018-Earnings-Release.pdf
24. Felt, A.P., et al.: Android permissions: user attention, comprehension, and behavior. In: Proceedings of the Eighth Symposium on Usable Privacy and Security (SOUPS 2012), Washington, D.C. ACM Press (2012). <https://doi.org/10.1145/2335356.2335360>
25. Fourniols, J.-Y., et al.: An overview of basics speech recognition and autonomous approach for smart home IOT low power devices. J. Signal Inf. Process. **9**, 239–257. <https://doi.org/10.4236/jsip.2018.94015>
26. de Freytas-Tamura, K.: The Bright-Eyed Talking Doll That Just Might Be a Spy (2018). <https://www.nytimes.com/2017/02/17/technology/cayla-talking-doll-hackers.html>
27. Fussell, S.: Behind Every Robot Is a Human (2019). <https://www.theatlantic.com/technology/archive/2019/04/amazon-workers-eavesdrop-amazon-echo-clips/587110/>
28. Ganjoo, S.: Is Facebook secretly listening your conversations? New report says yes, security experts say no proof (2018). <https://www.indiatoday.in/technology/features/story/is-facebook-secretly-listening-your-conversations-new-report-says-yes-security-experts-say-no-proof-1255870-2018-06-09>
29. Gao, G., Chow, M.: Android Applications, Can You Trust Google Play on These. Tufts University (2016)
30. Gharani, P., et al.: An Artificial Neural Network for Gait Analysis to Estimate Blood Alcohol Content Level. arXiv Comput. Res. Repos. abs/1712.01691 (2017). <https://arxiv.org/abs/1712.01691>
31. Google: Android 9 Pie. <https://www.android.com/versions/pie-9-0/>
32. Greenberg, A.: The Gyroscopes in Your Phone Could Let Apps Eavesdrop on Conversations (2014). <https://www.wired.com/2014/08/gyroscope-listening-hack/>
33. Grosche, P., et al.: Audio content-based music retrieval. In: Müller, M., et al. (eds.) Multimodal Music Processing. Dagstuhl Follow-Ups. Dagstuhl Publishing, Wadern (2012)

34. Hale, J.L.: Does Your Smartphone Listen To You? A New Study Debunked This Common Conspiracy (2018). <https://www.bustle.com/p/does-your-smartphone-listen-to-you-a-new-study-debunked-this-common-conspiracy-9682413>
35. Han, J., et al.: ACComplice: location inference using accelerometers on smartphones. In: 2012 Fourth International Conference on Communication Systems and Networks (COMSNETS), pp. 1–9 (2012). <https://doi.org/10.1109/COMSNETS.2012.6151305>
36. Han, J., et al.: PitchIn: eavesdropping via intelligible speech reconstruction using non-acoustic sensor fusion. In: Proceedings of the 16th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN), pp. 181–192. ACM Press, Pittsburgh (2017). <https://doi.org/10.1145/3055031.3055088>
37. Hashim, N.W., et al.: Evaluation of voice acoustics as predictors of clinical depression scores. *J. Voice* **31**(2), 256.e1–256.e6 (2017). <https://doi.org/10.1016/j.jvoice.2016.06.006>
38. Hassan, B.: 1 in 5 Aussies convinced their smartphone is spying on them (2018). <https://www.finder.com.au/press-release-july-2018-1-in-5-aussies-convinced-their-smartphone-is-spying-on-them>
39. He, Y., et al.: Dynamic privacy leakage analysis of Android third-party libraries. In: 1st International Conference on Data Intelligence and Security (ICDIS), pp. 275–280 (2018). <https://doi.org/10.1109/ICDIS.2018.00051>
40. Khatibloo, F.: Is Facebook Listening (And So What If They Are)? (2017). <https://www.forbes.com/sites/forrester/2017/03/17/is-facebook-listening-and-so-what-if-they-are/>
41. Kleinman, Z.: Is your smartphone listening to you? (2016). <https://www.bbc.com/news/technology-35639549>
42. Kröger, J.: Unexpected inferences from sensor data: a hidden privacy threat in the internet of things. In: Strous, L., Cerf, V.G. (eds.) Internet of Things. Information Processing in an Increasingly Connected World. IFIP Advances in Information and Communication Technology, vol. 548, pp. 147–159. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-15651-0_13
43. Kröger, J.L., et al.: Privacy implications of accelerometer data: a review of possible inferences. In: Proceedings of the 3rd International Conference on Cryptography, Security and Privacy (ICCSP). ACM, New York (2019). <https://doi.org/10.1145/3309074.3309076>
44. Lee, D.: Google admits error over hidden microphone (2019). <https://www.bbc.com/news/technology-47303077>
45. Liu, X., et al.: Discovering and understanding Android sensor usage behaviors with data flow analysis. *World Wide Web* **21**(1), 105–126 (2018). <https://doi.org/10.1007/s11280-017-0446-0>
46. Lomas, N.: Uber to end controversial post-trip tracking as part of privacy drive (2017). <http://social.techcrunch.com/2017/08/29/uber-to-end-controversial-post-trip-tracking-as-part-of-privacy-drive/>
47. Maheshwari, S.: That Game on Your Phone May Be Tracking What You’re Watching on TV (2017). <https://www.nytimes.com/2017/12/28/business/media/alphonso-app-tracking.html>
48. Mannini, A., et al.: Activity recognition using a single accelerometer placed at the wrist or ankle. *Med. Sci. Sports Exerc.* **45**(11), 2193–2203 (2013). <https://doi.org/10.1249/MSS.0b013e31829736d6>
49. Marczak, B., et al.: Hacking Team and the Targeting of Ethiopian Journalists (2014). <https://citizenlab.ca/2014/02/hacking-team-targeting-ethiopian-journalists/>
50. Marra, C.J., et al.: Ranking of News Feed in a Mobile Device Based on Local Signals (Pub. No.: US20170351675A1) (2017). <https://patents.google.com/patent/US20170351675A1/en>
51. Martínez, A.G.: Facebook’s Not Listening Through Your Phone. It Doesn’t Have To (2017). <https://www.wired.com/story/facebooks-listening-smartphone-microphone/>

52. McAfee: Net Losses: Estimating the Global Cost of Cybercrime. Center for Strategic and International Studies (CSIS), Washington, D.C. (2014)
53. McLaren, M., et al.: The 2016 speakers in the wild speaker recognition evaluation. In: Proceedings of the 16th Annual Conference of the International Speech Communication Association (INTERSPEECH), pp. 823–827 (2016). <https://doi.org/10.21437/Interspeech.2016-1137>
54. Michalevsky, Y., et al.: Gyrophone: recognizing speech from gyroscope signals. In: Proceedings of the 23rd USENIX Security Symposium, pp. 1053–1067 (2014)
55. Mohapatra, P., et al.: Energy-efficient, Accelerometer-based Hotword Detection to Launch a Voice-control System. (Patent No.: US20170316779A1) (2017). <https://patents.google.com/patent/US20170316779A1/en>
56. Morris, I.: Android Is Still Failing Where Apple’s iOS Is Winning (2018). <https://www.forbes.com/sites/ianmorris/2018/04/13/android-is-still-failing-where-apples-ios-is-winning/>
57. Naor, I.: Breaking The Weakest Link Of The Strongest Chain (2017). <https://securelist.com/breaking-the-weakest-link-of-the-strongest-chain/77562/>
58. Nichols, S., Morgans, J.: Your Phone Is Listening and it’s Not Paranoia (2018). https://www.vice.com/en_uk/article/wjbzzy/your-phone-is-listening-and-its-not-paranoia
59. Pan, E., et al.: Panoptispy: Characterizing Audio and Video Exfiltration from Android Applications. Proc. Priv. Enhanc. Technol. **2018**(4), 33–50 (2018). <https://doi.org/10.1515/popets-2018-0030>
60. Perlroth, N.: Governments Turn to Commercial Spyware to Intimidate Dissidents (2017). <https://www.nytimes.com/2016/05/30/technology/governments-turn-to-commercial-spyware-to-intimidate-dissidents.html>
61. Polzehl, T.: Personality in Speech. Springer, Cham (2015). <https://doi.org/10.1007/978-3-319-09516-5>
62. Quattrone, A.: Inferring Sensitive Information from Seemingly Innocuous Smartphone Data. The University of Melbourne (2016)
63. Rahman, M., et al.: Search rank fraud and malware detection in Google Play. IEEE Trans. Knowl. Data Eng. **29**(6), 1329–1342 (2017). <https://doi.org/10.1109/TKDE.2017.2667658>
64. Ramirez, E., et al.: Data Brokers. A Call for Transparency and Accountability. Federal Trade Commission, Washington, D.C. (2014)
65. Ramirez, R., et al.: Cross-Device Tracking: An FTC Staff Report. Federal Trade Commission, Washington, D.C. (2017)
66. Rosenbach, M., et al.: iSpy: How the NSA Accesses Smartphone Data (2013). <http://www.spiegel.de/international/world/how-the-nsa-spies-on-smartphones-including-the-blackberry-a-921161.html>
67. Schlegel, R., et al.: Soundcomber: a stealthy and context-aware sound trojan for smartphones. In: Proceedings of the Network and Distributed System Security Symposium (NDSS) (2011)
68. Schmidt, D.C.: Google Data Collection. Digital Content Next, New York (2018)
69. Sidor, S.: Exploring limits of covert data collection on Android: apps can take photos with your phone without you knowing (2014). <http://www.ez.ai/2014/05/exploring-limits-of-covert-data.html>
70. Statista: Global mobile OS market share in sales to end users from 1st quarter 2009 to 2nd quarter 2018. <https://www.statista.com/statistics/266136/global-market-share-held-by-smartphone-operating-systems/>
71. Stern, J.: Facebook Really Is Spying on You, Just Not Through Your Phone’s Mic (2018). <https://www.wsj.com/articles/facebook-really-is-spying-on-you-just-not-through-your-phones-mic-1520448644>

72. Tang, Q., et al.: Automated detection of puffing and smoking with wrist accelerometers. In: Proceedings of the 8th International Conference on Pervasive Computing Technologies for Healthcare. pp. 80–87 (2014)
73. Taylor, P.: Edward Snowden interview: “Smartphones can be taken over” (2015). <https://www.bbc.com/news/uk-34444233>
74. 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). <https://doi.org/10.1145/2750858.2807545>
75. Timberg, C., et al.: WikiLeaks: The CIA is using popular TVs, smartphones and cars to spy on their owners (2017). https://www.washingtonpost.com/news/the-switch/wp/2017/03/07/why-the-cia-is-using-your-tvs-smartphones-and-cars-for-spying/?noredirect=on&utm_term=.c162373021c3
76. Triggs, R.: No, your smartphone is not always listening to you (2018). <https://www.androidauthority.com/your-phone-is-not-listening-to-you-884028/>
77. Tsukayama, H., Romm, T.: Lawmakers press Apple and Google to explain how they track and listen to users (2018). <https://www.washingtonpost.com/technology/2018/07/09/lawmakers-press-apple-google-explain-how-they-track-listen-users/>
78. Yerukhimovich, A., et al.: Can smartphones and privacy coexist? Assessing technologies and regulations protecting personal data on Android and iOS devices. MIT Lincoln Laboratory, Lexington, MA (2016). <https://doi.org/10.7249/RR1393>
79. Zhang, L., et al.: AccelWord: energy efficient hotword detection through accelerometer. In: Proceedings of the 13th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys), pp. 301–315. ACM Press (2015). <https://doi.org/10.1145/2742647.2742658>
80. No, Phones Aren’t Listening to Your Conversations, but May Be Recording In-App Videos: Study (2018). <https://www.justandroid.net/2018/07/05/no-phones-arent-listening-to-your-conversations-but-may-be-recording-in-app-videos-study/>

Open Access This chapter is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, duplication, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, a link is provided to the Creative Commons license and any changes made are indicated.

The images or other third party material in this chapter are included in the work’s Creative Commons license, unless indicated otherwise in the credit line; if such material is not included in the work’s Creative Commons license and the respective action is not permitted by statutory regulation, users will need to obtain permission from the license holder to duplicate, adapt or reproduce the material.

