

Chapter 11

Privacy and the Internet of Things



**Heather Richter Lipford, Madiha Tabassum, Paritosh Bahirat, Yaxing Yao,
and Bart P. Knijnenburg**

Abstract Using networks of Internet-connected sensors, the Internet of Things (IoT) makes technologies “smart” by enabling automation, personalization, and remote control. At the same time, IoT technologies introduce challenging privacy issues that may frustrate their widespread adoption. This chapter addresses the privacy challenges of IoT technologies from a user-centered perspective and demonstrates these prevalent issues in the domains of wearables (e.g., fitness trackers), household technologies (e.g., smart voice assistants), and devices that exist in the public domain (e.g., security cameras). The chapter ends with a comprehensive list of solutions and guidelines that can help researchers and practitioners introduce usable privacy to the domain of IoT.

11.1 Defining IoT

The Internet of Things (IoT) is revolutionizing our use of computing, introducing networked devices throughout our everyday lives that collect and utilize information to provide an ever-growing number of services. Coined by Kevin Ashton during a presentation at Proctor and Gamble, the Internet of Things primarily originated

H. R. Lipford (✉) · M. Tabassum
College of Computing and Informatics, UNC Charlotte, Charlotte, NC, USA
e-mail: richter@uncc.edu; mtabassu@uncc.edu

P. Bahirat · B. P. Knijnenburg
School of Computing, Clemson University, Clemson, SC, USA
e-mail: pbahira@clemson.edu; bartk@clemson.edu

Y. Yao
Department of Information Systems, University of Maryland Baltimore County, Baltimore, MD,
USA
e-mail: yaxingya@umbc.edu

with the idea of RFID tags to be used for the purpose of streamlining supply chain operation [1]. A broad definition of IoT is:

The Internet of Things refers to the unique identification and ‘Internetization’ of everyday objects. This allows for human interaction and control of these ‘things’ from anywhere in the world, as well as device-to-device interaction without the need for human involvement

While IoT was originally conceptualized for industrial and manufacturing domains, the concept has found its place in numerous areas, ranging from public domains such as smart cities to the most intimate parts of our lives with smart homes and fitness trackers. From energy and health monitoring to remote operation and surveillance, IoT devices provide exciting services to improve our lives [2]. Yet these devices also bring unique privacy challenges due to their integration into the world around us, and the extensive amount of data that they can collect and use. We first introduce the various domains of IoT, to summarize the kinds of data they collect and their uses before delving into the challenging privacy issues that the Internet of Things raises.

Broadly speaking, there are three core domains that fit under the umbrella of IoT. They are:

- **Wearable IoT**—devices that people can wear as accessories, such as watches, for monitoring an individual’s activities or vital signs.
- **Household IoT**—devices that sit in people’s homes, such as smart speakers, appliances, and thermostats.
- **Public IoT**—devices that are used in public places, such as smart water meters, autonomous vehicles, and Bluetooth beacons.

11.1.1 Wearable IoT Domain

IoT-enabled wearables are internet-connected devices integrated with various sensors that can be worn as external accessories (i.e., watches, glasses, rings, etc.) or implanted in textiles such as smart shoes or jackets. Many commercially available wearables are specifically targeted for health and fitness monitoring, such as Fitbit and Apple Watch. Sensors collect the movement and vital signs of individuals, such as steps taken, heart rate, and sleep quality, in order to track activities and to help people monitor and improve their wellness and physical performance [3]. Other devices aim to help users monitor their interaction with the world around them, such as Google Glass, allowing people to capture audio and video of their daily lives. All wearable devices share a common goal of automatically and unobtrusively recording an individual’s physical interaction with the world.

Despite the fact that much of this information can be related to an individual’s health, research suggests that users are comfortable sharing their information with a range of other people and organizations to support their health goals because they perceive much of that information, such as step count, as not particularly sensitive [4]

11.1.2 Household IoT Domain

A smart home refers to a residence that has lighting, heating, air-conditioning, security systems, or entertainment systems which communicate with one another and work together to improve the experience and increase the comfort of the occupants. Smart home devices allow for remote monitoring and operation of parts of a home, such as the thermostat, lights, or door locks. Many smart devices aim to increase the convenience and automation of the home, such as with smart speakers and appliances. Devices can also enable safety and security monitoring, using cameras, audio, and fire or water leak sensors. Additional people and organizations may also be involved in safety monitoring, with information and devices being shared with family members, security companies, or emergency services.

The perception of the privacy of smart home data varies by device. Some information is not perceived as very sensitive, such as the status of smart lights or thermostats. Yet, video and audio from inside the home are usually considered private and users desire strong protections against recordings being accessed without their knowledge or control [5, 6].

11.1.3 Public IoT Domain

IoT technology has also gained popularity in public infrastructure through smart cities and smart buildings. Public IoT infrastructure brings a number of benefits in the management and optimization of traditional public services, such as transportation and parking, lighting, ventilation, surveillance and maintenance of public areas, and even preservation of cultural heritage. For example, the New York City Department of Transportation integrated a congestion management system to determine traffic speed at 23 intersections in Midtown Manhattan that has improved the travel time by 10% in Midtown's avenues [7]. Similar to smart homes, smart cities and buildings also provide services to monitor the security and safety of spaces and people and intelligently automate controls in response to the environment. In addition, IoT is frequently used for resource management, lowering costs by more efficiently and intelligently utilizing resources. For example, the city of Dallas, Georgia, has undertaken a smart water meter program, which helped them to detect water leaks more efficiently and minimize water loss [8].

Another emerging type of IoT device in the public domain is the autonomous vehicle, which is increasingly adopted in app-based taxi services (e.g., Uber), home delivery services, and consumer products (e.g., Tesla). Each autonomous vehicle is equipped with a large amount of sensors to collect information about the surrounding environment, including people who are walking on the street, other vehicles on the road, and nearby store information [9]. In addition, the drivers and passengers who sit in the car also face a large amount of data collection during and after their ride (e.g., vehicles may collect information about their daily schedule) [10].

11.1.4 Outline

In the following sections, we will further discuss the unique privacy challenges introduced by the use of IoT devices. To further illustrate these challenges in practice, we then discuss them in more detail through three case studies of fitness trackers, smart home voice assistants, and CCTV and surveillance cameras. While research into solutions to these challenges is still limited, we end the chapter with a discussion of potential ways to reduce the privacy risks users face and address user needs in managing their privacy with IoT.

11.2 Privacy Challenges

Similar to other technologies, IoT devices face many types of privacy issues. However, due to the volume of data collected, the ubiquitous nature of IoT devices, and their ability to blend into the background, they also introduce new challenges and greatly exacerbate existing privacy challenges when compared with traditional computing applications. In this section, we highlight the key challenges, which include the following:

- People **lack awareness** of the data practices of IoT devices and their manufacturers, due to the large amount of data involved in potentially complex ecosystems of devices, as well as the unobtrusive methods of data collection.
- The **accumulation** of large amounts of data enables the **inference** of sensitive information, unbeknownst to users.
- IoT devices can be used by **multiple users** and in environments containing **multiple other people**, increasing the complexity of privacy needs and access controls.
- IoT devices offer **limited controls** for users to manage the privacy of themselves and their information. Many scenarios and domains involve **bystanders**, who currently have **no ability to control** devices whatsoever.
- **Security mechanisms and processes** to protect the devices and data collected may not be robust or mature, putting users' privacy at risk due to attacks and data breaches.

11.2.1 No Awareness/No Interface

The decisions people make regarding their usage of a given computing device is governed by their *mental model* of the device, which is comprised of their assumptions and intuitions regarding what data they think is collected, how that data is used, how it is stored and shared, etc. This awareness is primarily based on the experiences people have with their devices over time—the kinds of exchanges

they have with an application and the information involved in those exchanges. IoT devices differ from traditional computing and mobile devices in that they are more embedded into the surrounding environment, often without a dedicated screen, resembling non-computing devices and yet unobtrusively capturing and utilizing a range of information.

For instance, the “Hello Barbie” doll looks like a typical kids’ toy, “Alexa” functions as a music speaker, and fitness trackers usually resemble traditional watches. Data collection is mostly invisible and automatic. Beyond the devices owned by the user, any environment they enter may potentially have devices owned by others, each collecting their own unique set of information. Together, this means that users cannot rely on their former perceptions of what interacting with a doll, a speaker, or a watch means. Users may also become habituated to their devices, and hence gradually become less aware of the pervasive data collection. In other words, users must form new mental models of IoT devices, and as these mental models are based on the incomplete information they receive from their interaction with the devices, it is not surprising that they tend to make incorrect assumptions about the privacy of their IoT devices (e.g., they may think that their child’s “Hello Barbie” doll does not collect and store any data). These incorrect assumptions can lead to privacy intrusions. Alternatively, users may choose to not adopt a given device over privacy fears that arise from being uncertain about the data collection practices of the device. In this case, their uncertainty leads them to forego the adoption of technology that they would otherwise be comfortable with.

Studies of current smart home users have demonstrated that people are generally aware of the collection and use of information that is apparent in the functioning of the device [11]. For example, users understand that a smart thermostat captures temperature changes, a security camera records video, and a smart lock logs when the door is locked and unlocked. Users expect that this collected information is used by the device to properly function and provide useful services, and potentially by the manufacturers to improve their devices. Yet, while users do expect this information to be stored in the “cloud,” and not on the device itself, there is little awareness over exactly what that means [11]. Users are unclear on how, where and for how long their information is stored, who it could be shared with, and what other uses could occur [11, 12]. Studies have demonstrated similar perceptions for wearable devices [13, 14]. Yet, few have studied the perceptions of IoT devices in more public settings, where users likely have less awareness of the presence of devices and little interaction with them.

The standard method for users to know about what data is collected and how it is used is the privacy policy or the end-user license agreement provided by vendors. These methods are already problematic for traditional computing devices, with few people reading them. Yet they are even more difficult to rely on for IoT devices. Studies have shown that the boxes and print materials of smart home devices rarely describe the device’s data collection practices [15, 16]. As devices themselves have a small screen or no screen at all, users must instead visit a separate website or use an accompanying app to view the policy. Even if someone actively looks for information on a vendor’s website, the privacy policy may provide information only

about the website data practices, not the data practices regarding the device's sensor data [15]. Furthermore, the only people likely to view any sort of policy statement are those doing the setup and installation of the device. Others who are in the purview of the device (e.g., other home or building occupants) will not have this opportunity. This is particularly the case for IoT devices in public spaces, where those whose data is collected may not be aware of the existence of the device at all. As such, providing transparency regarding the ownership and policies of the data collected by IoT devices in smart buildings or public infrastructure is even more challenging.

A final challenge is the complex ecosystem in which many IoT devices are embedded. Not only the device itself exchanges information with the manufacturer; users often interact with an accompanying mobile app to access and control the device and its information, and this app itself may perform additional data collection, such as tracking the location of the user. Furthermore, devices may be interconnected with—and share data with—smart hubs and other devices, which may be built by a different manufacturer. Finally, third-party applications may operate on top of any of these platforms and involve an additional exchange of information between organizations. Even for tech-savvy users it is very difficult to fully understand how information is collected, stored, and shared by each of these entities. Yet, many users have a fairly simple service-oriented view of how different devices interact. For example, a smart home user may know that they can turn on their TV using their Google Home device, but they will have little knowledge of what information is exchanged between the TV and Google to accomplish this task [11, 17].

We summarize these challenges as follows:

- **Devices are unobtrusive** and often do not “look” like they are collecting extensive amounts of data.
- **Users do not understand the extent of data collection** and how data may be used for secondary purposes.
- **Users do not read the privacy policy** or may simply not have access to it.
- **The IoT ecosystem is complex** and understanding, let alone managing, the data collection practices of multiple actors is a huge undertaking.

11.2.2 Accumulation and Inference of Data

Another unique issue with IoT devices is the sheer volume of data that is collected, from so many different sources. Taken by itself, each individual piece of data seems innocuous—step counts or the status of a light in a house are not considered sensitive [4]. Yet over time, this accumulation of information can allow applications to learn powerful patterns of human preferences and behaviors. For example, the data from wrist-worn IoT devices such as smartwatches and fitness trackers can be used to infer users' physical activities such as walking, running, and jumping, with high

accuracy [18]. While this may be somewhat expected, wrist-worn devices can even allow inferences about what the user is typing [19]. Similarly, from smart meter data, it is possible to recognize bathroom activities, cooking, and housework [20].

The threat of profiling increases when the large amount of data collected by IoT devices is *aggregated* to reveal previously inaccessible parts of people's lives. The aggregation of data from multiple devices can provide sensitive information about the users that cannot be determined from an individual data source alone. For instance, a thermostat with temperature zone control knows about the collective movement of occupants in the house. When users control the thermostat using a smartphone, then the thermostat can learn exactly who is where in the home and when. Inferences based on data aggregated from multiple devices of data can cause an unexpected revelation of users' identity, personal traits, activities, habits, preferences, sexual orientation, health status, financial situation, and more, even when data is collected anonymously [15]. A system's ability to make such inferences is beyond most users' comprehension. Indeed, even if users have some idea about the data collected by each individual device, they will likely be unable to understand the privacy implications of the aggregated data from multiple devices.

Information can even be inferred from the metadata and communication patterns of devices without gaining access to the data itself. For example, network traffic rates from a Sense sleep monitor reveal consumer sleep patterns, network traffic rates from a Belkin WeMo switch reveal when a physical appliance in a smart home is turned on or off, and network traffic rates from a Nest Cam Indoor security camera reveal when a user is actively monitoring the camera feed or when the camera detects motion in a user's home [21]. This is alarming, as Internet Service Providers (ISP) have easy access to traffic data, and the US legislature voted in 2017 to allow ISPs to use and sell the data collected from their customers' network traffic [22].

Several studies have examined users' expectations and concerns when it comes to IoT inferences. Studies of wearable devices reveal that users are concerned about sharing data with insurance companies, for example, as the information could be used to raise rates [4, 23]. While users have some expectations that their behaviors and habits could be inferred, they are unsure and lack awareness of the kinds of scenarios that are already plausible [23, 24].

The most immediate and obvious use of aggregated data is to create a reasonably accurate profile of a user for the purpose of advertising. For instance, Amazon and Google have patented the use of digital voice assistants to extract keywords from ambient speech and to use those keywords to provide relevant advertisements [25]. Studies of IoT inferences have shown that based on their experiences with web browsers and Internet applications, users do have expectations that organizations will use their information to target advertisements [11, 12]. Users are not overly concerned with such advertisements, even though they can at times be creepy. However, users are concerned that such information could also be used to manipulate their behavior; for instance, users can be influenced to buy a certain product they do not want or nudged to spend more money [26].

There are few mechanisms to educate users about the potential implications of inferred information, and few concrete threats have yet been reported. Yet,

we have already seen concern and even backlash over profiling and inferences in other domains, particularly social media. For example, several scandals, such as the Cambridge Analytica scandal, have left millions of Facebook users surprised and dismayed over the use of their information to create political profiles or infer their moods outside of their awareness [27–29]. Such examples will only increase as organizations figure out how to capitalize on their IoT data.

We summarize the data accumulation and inference challenges as follows:

- **The sheer volume of IoT data threatens user privacy** as it can be used to infer users' private activities.
- **The aggregation of data across multiple devices** further increases the threat and is much harder for an end-user to comprehend.
- **Even metadata can reveal sensitive information** and such data is available to users' ISPs who are allowed to sell it.
- **Unwanted inferences are likely to generate backlash** especially if they go beyond targeted advertising (e.g., the Cambridge Analytica scandal).

11.2.3 Multiple Users

IoT devices are used by and around a variety of people. This is particularly salient for IoT devices that support smart building and city infrastructure: these devices are intended to track activities and behaviors of potentially large numbers of people. For instance, the city of San Diego has cameras built into its streetlights, which capture pedestrian traffic [30]. Likewise, driving patterns can be captured by the connected cars program piloted by New York City [31]. Even in a household setting, IoT devices capture data of a large number of people, including the multiple family members or roommates who live in the home, family and friends who visit, and house cleaners and contractors who help with maintenance. A user may also share remote access to their smart home devices with people outside of their home. For instance, neighbors could check on each other's homes in case of a fire or burglar alarm, or share access to each other's security or doorbell cameras to monitor community safety and security. Friends or family members could remotely check on pets, or let in people delivering packages, should the homeowner not be available [11]. And while wearable and health devices are primarily designed for single users, they are also commonly shared among different household members to gain their benefits without the expense of additional devices [32]. Wearable users may also share information with caregivers or doctors to receive timely medical intervention [23].

One critical privacy implication of this multi-user environment is that users may have complex preferences for how to share access to and control over their IoT devices and the collected data with others, especially if they have different social relationships with those others. For instance, if IoT devices are shared between immediate family members such as a spouse or adult children, users will likely be

comfortable sharing sensitive controls (i.e., the ability to order something through the smart speaker) and information (i.e., the ability to see health profiles in fitness trackers) because of the high level of trust. However, when sharing with less trusted users (i.e., visitors, roommates, neighbors, house help) or under-aged users (i.e., kids, teenagers), people may have more restrictive access control preferences based on device capabilities and other contexts [11, 33]. For instance, home owners tend to be more comfortable sharing the live view feature of an outdoor camera with neighbors than with sharing the same capability of an indoor camera [33]. However, users may want to share that same indoor camera with neighbors when they are out of town or in case of an alarm [11].

The complexity of IoT users' access control needs can be addressed by time-based access control to share temporary access (e.g., a one-time key to drop off a package), location-based access control to share access based on the location of the user (e.g., monitoring the house when the user is away), role-based access control to grant or restrict certain capabilities for certain user roles (e.g., to prevent young children from ordering products via a smart speaker), and event-based access control to share only specific capabilities required for a particular event (e.g., to alert emergency services in case of an alarm). However, current devices are very limited in the kinds of controls they provide. Furthermore, current controls make it difficult to understand what access rights are being shared [11]. Hence, there is a risk of oversharing sensitive information (e.g., video recording of household members) or control (e.g., allowing the deletion of video recordings) with other users. This lack of transparency and existence of adequate access control mechanisms leads users to share everything with their most trusted community, often by sharing full account credentials, and to not share the device at all with people who are less trusted [11]. Yet while additional controls may enable more fine-grained access, they run the risk of introducing too much complexity, which may overwhelm users, leading instead to even more loss of user control. Hence, the challenge is to understand the most prevalent sharing scenarios and needs for different devices and platforms to decide (1) how to prioritize between different access control mechanisms for different devices and (2) how to balance users' complex access sharing requirements with their need to share devices without much effort.

Even for devices shared freely between multiple people, different users may have different preferences regarding what information and capabilities they find sensitive and how information should be shared in different contexts. To add to that complexity, users may have a different level of interaction and control of the devices. For instance, in a smart home context, admin users who set up and maintain the devices have more control and power and may be able to violate the general expectations of privacy of others who have limited control over the devices [32, 34]. Resolving users' conflicting preferences regarding the use and control of shared devices remains a challenge, as does the prioritization of the privacy needs of users in different roles.

The challenges of multi-user IoT devices can be summarized as follows:

- **IoT devices are regularly used by multiple users**, even those devices that are designed with a single user in mind.
- **IoT users tend to have complex sharing preferences** that depend on their relationship with the other users.
- **Many IoT devices lack the mechanisms to support these preferences** and make it difficult to understand what is being shared with whom.
- **Reconciling the privacy needs of different users remains a challenge** that is not adequately addressed by existing IoT devices.

11.2.4 Little Control

The greatest strength of IoT lies in automation: IoT devices can take over routine or mundane tasks that would otherwise be performed by humans, thereby providing convenience. This means that to realize the full benefits of IoT, users must relinquish some level of control [35]. The tension between automation and the need for control becomes even more important in case of privacy.

A major issue with control revolves around ownership: those who interact with or are subject to the data collection practices of a device may not necessarily be the owners of the device. In addition, the owners of a device may not be the owners of the data collected by the device. Indeed, scholars studying the ownership of IoT systems have called IoT an “Imminent Ownership Threat” [36]. Their concerns revolve around questions of who has control over a smart device’s actions, as well as who owns and manages the data collected by the device [37]. The latter is particularly complicated in the case of shared IoT systems. Take the case of an Airbnb host, who technically owns the smart devices installed on the rental property. Renters of the property are likely concerned about their privacy, especially in case of surveillance cameras, and may assert that they should have control over these devices for the duration of their stay. At the same time, the host would like to maintain control over their devices, for example, to ensure the safety of the property [38].

The complexity of such tensions between parties regarding the ownership of recorded data is further exaggerated in public IoT systems, mainly due to the increased number of parties who are subject to the devices’ data collection practices [36, 37]. Indeed, one of the most challenging aspects of IoT is the involvement of bystanders who have no control over—or in many cases even awareness of—the devices that collect data about them and the capabilities of these devices [6]. Most IoT devices leave few opportunities for bystanders to be notified of, or give consent to, being recorded by devices in their surroundings, such as when being captured by a neighbor’s smart doorbell as one walks down the sidewalk. In most cases, the only preferences a bystander may be able to express are the basic decision of whether or not to enter a space. While this is clearly an issue for smart cities,

these issues can still occur with serious consequences in more intimate settings. An example is the considerable backlash over the deployment of Google Glass, as it was difficult for a bystander to determine whether they were being video-recorded by the person wearing the glasses, making those around a user feel uncomfortable [39]. While researchers have investigated a few potential technological solutions for bystanders, such as automatically obfuscating faces in videos of bystanders, there are few mechanisms or policies currently deployed to reduce these tensions between device and data owners, and the many additional people who are captured and impacted by those devices.

Even when users do have ownership of the device and access to its privacy controls, these controls are often quite limited. For example, as we discussed in Sect. 11.2.3, users may have complex needs for controlling the amount of access others have to their smart home devices. Yet, access control capabilities are often so limited, or lack transparency as to what they allow, that users provide access by simply sharing full account credentials with only their closest family and friends [11]. Other studies have shown that users often do not use existing privacy control mechanisms [40], such as the ability to review and delete recorded conversations, and may not even be aware of such mechanisms. Manufacturers of consumer IoT devices have tried to cater to this issue by providing privacy mechanisms that physically situated on the device, for example, Google Nest and Echo Dot each come with a physical button to disable microphones [41], and Facebook's Portal devices come with an integrated camera shutter or with physical camera covers [42]. These features increase the visibility of privacy capabilities and give users confidence that the mechanism is actually performing as intended. However, such physical privacy features are naturally limited in their complexity.

We summarize the challenges surrounding the lack of control in IoT systems as follows:

- **IoT systems create complex issues around ownership and control** and must find intuitive ways to address those issues.
- **IoT systems may violate the privacy of bystanders** and give them little opportunity to become aware of, let alone take control over, the collection practices that they are subjected to.
- **IoT privacy controls are often limited even for main users** leading to suboptimal privacy management practices.
- **Physical privacy controls can raise trust and awareness** but are often limited by their rudimentary functionality.

11.2.5 (In)Security of IoT Devices

IoT devices create a large number of attack vectors, resulting in many possibilities for adversaries to compromise the devices and use them for nefarious purposes. Successful attacks then compromise the privacy of device users and their information.

For example, in recent news, we saw a number of successful security attacks on smart home devices, such as the Mirai botnet (a DDOS attack on networked devices running Linux) [43], the monitoring of home occupants via their thermostat [44], the unauthorized access to google calendar information from a smart fridge [45], and the compromising of baby monitors to allow external parties to monitor live feeds, change the camera settings, and authorize other users to remotely view and control the device [46, 47]. Smart home devices, in particular, are becoming an easy and lucrative target for malicious attackers because of the availability of insecure devices and the fact that compromising one device can allow them to compromise several other connected devices in the same network.

Given that security is such a critical issue in IoT devices, it is remarkable that many devices do not have appropriate security mitigations in place. This issue has multiple root causes. The first root cause of this problem is that most IoT devices are connected directly to the Internet, which exposes them to all the network security problems of a typical online system. IoT device networks are extremely heterogeneous; they can consist of a large number of different devices, applications, and communication technologies. As such, there is not one universal security solution that can decrease or mitigate all of the security risks for all of these devices. Moreover, IoT devices typically do not have enough processing resources to support traditional security mechanisms [48]. Finally, the inter-connected nature of IoT devices contributes to their vulnerability, because even though some devices may have relatively strong security mechanisms, their security can still be compromised through other, less secure devices that they are connected to.

A second root cause is that manufacturers do not focus enough on security when developing their products—particularly for consumer-oriented devices such as those found in smart homes. Due to the novelty of the domain, many IoT devices are developed as quickly and cost-effectively as possible in an attempt to compete in the already crowded market place [49]. Security requirements are likely to take a lower priority than other features and functionality, and with less awareness of the risks consumers may not demand or pay for additional security protections. Thus, many IoT devices do not implement common security mechanisms such as encrypted communication, making them vulnerable to security attacks [50]. Many manufacturers have yet to establish mature security processes and have not yet allocated the resources needed to invest in substantial vulnerability detection and mitigation. When security is initially ignored in traditional software applications, these vulnerabilities are usually fixed over time through updates and security patches. However, not all IoT devices allow for regular and automated software updates to patch vulnerabilities [51].

Finally, another factor that makes IoT devices vulnerable to cybersecurity attacks are the users themselves. Most users of wearable and smart home devices who set up those devices are not professionals. They may not know the security risks imposed by networked devices and how to protect their devices against those risks. Indeed, some users may erroneously believe that traditional security practices, such as using strong passwords, are enough to protect them against security risks in their homes, as they do not fully understand the threats inflicted on them by their smart

devices [48]. Even if users do know about security measures, these measures tend to be too complicated for them to implement correctly [11]. Finally, IoT researchers and manufacturers have not yet developed clear guidelines and best practices for users to help them employ appropriate practices, reducing their risks of security attacks [11, 17].

In summary, the security vulnerabilities of IoT devices fall along the following lines:

- **IoT devices introduce a significant security threat**, and it is therefore remarkable that most consumer-facing IoT devices lack proper security protections.
- **The heterogeneous and Internet-connected nature of IoT systems makes them difficult to secure**, and vulnerabilities in one device may leave other devices in the network vulnerable as well.
- **Market pressures and limited device capabilities make it difficult to provide proper security**, and security patches may take a long time to propagate within the network.
- **Users may not be capable of setting up their devices in a secure manner** and may not fully understand the threats caused by their IoT devices.

11.3 Case Studies

In this section we discuss three case studies and illustrate the privacy issues that have arisen in these cases. In particular, we spotlight the following three IoT devices: wearable fitness trackers, household smart voice assistants, and CCTV and smart cameras. To summarize this section:

- In the **wearable domain**, **fitness trackers** collect data that is largely considered nonsensitive, and users share their data with a variety of other people and organizations to help meet their fitness and health goals. However, users' lack of awareness of potential health-related inferences are considered more sensitive.
- In the **household domain**, **smart voice assistants** collect audio, which can be viewed as intrusive, despite controls and features that limit that collection.
- In the **public domain**, **security cameras** cause people to change their behavior when they perceive they are being watched. In public, CCTVs can result in less anti-social behavior and reduce crime. Yet, as smart cameras move into more private spaces, constantly being watched may have a chilling effect on behavior, particularly for those who lack control over the cameras.

11.3.1 Fitness Trackers

Fitness trackers are wearable devices that have gained significant popularity in recent years, with brands including Fitbit, Garmin, and Polar. Many take the form

of a wrist-worn device, but they can come in a range of form factors depending on the sensors and intended usage. Apple also has fitness tracking built into the Apple watch and iPhone. The primary use of a fitness tracker is to monitor different aspects of a user's health and fitness, to motivate a more active lifestyle, to track performance, or to monitor a health condition.

A unique aspect of fitness tracker usage is that users regularly share data captured by their device with other people for a variety of purposes [52]. Users seek accountability and mutual support for their health goals by sharing their fitness progress on social media or within community forums [53]. They share health-related data with medical providers or caregivers. Employers and insurance companies may incentivize data sharing in an effort to encourage healthy behaviors through fitness campaigns [23]. Thus, users face significant challenges in managing not just the collection of their information but the sharing and use of that information with a potentially large number of other people and organizations. While most devices offer a range of sharing controls, these controls do not always provide fine-grained customization of data sharing [54].

Researchers have also demonstrated a large number of inferences that can be made with fitness tracker data. Mood, stress level, places, and sexual activity can be determined with high accuracy [18, 55, 56]. For example, in January 2018, reports revealed that fitness tracker data shared by users on Strava, a social fitness service, showed accurate locations of US military sites [57]. Despite this potential, users seem generally unconcerned about the risks of sharing their information in such a public manner. For example, research shows that users do not consider sharing one's step count with a pharmacy a cause of privacy concern. Instead, users are more concerned about managing others' impressions of them and sharing information that fits the norms of various platforms [23]. For example, someone might not want to share their lack of exercise on social media lest friends might view them as lazy.

One reason for users' high level of comfort with sharing fitness tracking data is their lack of awareness of the possible inferences that can be made with such data. Studies have demonstrated that users do not believe that certain inferences are even possible or sufficiently accurate to be useful [58, 59]. This lack of awareness may be due to a dearth in application features that can inform users about the way their information could be aggregated and used to make various inferences. Thus, users may currently be comfortable sharing their step count or heart rate but may consider it a privacy invasion if they knew that more sensitive information about their health or activity has been inferred based on this data.

Providing privacy awareness and controls for fitness trackers is challenging, though: Their often tiny screens are barely large enough to fit the necessary functional information (e.g., time, heart rate, reminders, etc.), and hence opt to leave out all other information (e.g., what data is collected, how data is handled, etc.). Even though users can often access such information, as well as some controls, through the associated apps on their smartphones (e.g., device settings) or through a corresponding web portal (e.g., privacy policy), this decoupled way of interaction makes it less appealing for users and reduces opportunities to learn about data practices.

We summarize the privacy challenges of fitness trackers as follows:

- **Users are generally comfortable sharing their data** with friends, caregivers, and sometimes even employers and insurance companies.
- **Tracker data enables a large number of inferences** that revolve sensitive information, even if the underlying data itself is not regarded as sensitive.
- **Users lack awareness of the possible inferences that can be made**, which may explain their current openness to data sharing.
- **Giving users fine-grained control is challenging** given the small form factor of most fitness trackers.

11.3.2 *Smart Voice Assistants*

Voice commands have become one of the most prominent modes of interacting with smart technology, particularly in smart homes. Triggered by voice commands like “Hey Google” or “Alexa,” these assistants will listen to users’ questions or requests. Hence, these devices continuously listen for audio cues from their surroundings to respond the moment they are called on. In response to user queries these devices can provide audio feedback and carry out a variety of actions, both virtual and physical. For instance, users can buy something from Amazon through the Amazon Echo and receive notification of packages delivered. Moreover, smart voice assistants are often connected to and used to control other smart home devices. For instance, users can use voice commands to ask their smart voice assistant to turn on their smart lights or TV.

As smart assistants are increasingly embedded in everyday conversational settings, concerns have been raised by several researchers and journalists around the devices’ intrusive listening practices [40, 60]. There are general suspicions and confusion surrounding what is exactly being recorded by these devices and how the parent company handles the audio recordings. Several incidents of Amazon Echo sending sensitive recordings to someone without the owner’s knowledge and approval have been in the news [61], contributing to consumer concerns. Indeed, the intrusiveness of smart assistants, their potential to violate users’ privacy, and distrust of the companies that manufacture them are the main reasons reported for not adopting such devices [40, 62]. Although some companies proactively provide a set of privacy controls for smart assistants, end users are often unaware of these controls. For instance, a recent study found that most end users are not aware of their ability to view and delete the audio logs, even though those same users were not comfortable with the permanent retention of their recordings [63]. Moreover, some of the privacy controls are misaligned with users’ needs. For instance, Google Home and Amazon Echo offer a physical mute button that requires different interactions than regular voice commands, and hence the button is rarely used [40].

In addition to the concern over intrusive data collection practices, smart assistant owners also face the challenge of limiting others’ access to sensitive information and

actions that can be performed with the device, such as buying items. For example, in Texas, a 6-year-old was able to order a dollhouse and four pounds of cookies using Amazon Echo [64]. While users can add a voice code that must be used during shopping as an extra layer of protection, many users would not think to search for this capability, and even if used, the code must be spoken aloud and can easily be overheard.

These concerns and feeling of intrusiveness may heighten as smart assistants are finding their way into our cars. There is a clear benefit to enabling car owners to control different activities in the car via simple voice commands, allowing drivers to keep their hands on the wheel at all times. Consequently, a large number of insurance providers are giving away a car-based Alexa assistant for free to their users, citing the benefit of reducing accidents caused due to texting while driving [65]. However, this trend presents privacy challenges for car passengers, particularly in the case of ride-sharing scenarios such as Uber or Lyft.

With the increasingly seamless integration of smart assistants into our daily lives, they are likely to become even more intrusive. For example, Amazon and Google have both patented mechanisms for using their digital voice assistant to extract keywords from ambient speech provide targeted advertisements [66]. In the future, a voice assistant may proactively provide assistance based on users' conversation without being invoked by the wake word [67]. Such a proactive device has a tremendous potential for helping users by providing more personalized and contextual services [35]. However, the intrusiveness of such a device calls for extensive research to identify privacy features that would allow users to enjoy these benefits without having to worry about their privacy.

The privacy challenges of smart voice assistants can be summarized as follows:

- **Smart voice assistants proactively listen for audio cues**—an intrusion that causes many to avoid adopting them.
- **Users are often not aware of existing privacy controls**, as they tend to be “hidden in plain sight.”
- **Car-based voice assistants** can improve driver safety but are also intruding upon the privacy of passengers.
- **Future proactive voice assistants** have a tremendous potential to provide personalized services while at the same time further exacerbating users' privacy concerns.

11.3.3 Security Cameras

The “Watching Eye Effect” refers to the behavior modification that can occur upon the perception of being observed by something. Researchers have shown that this phenomenon can play an important role in reducing antisocial behavior of individuals in public [68, 69]. One could argue that such behavioral modification is an unwanted intrusion into people's lives, though. Moreover, the widespread

deployment of smart cameras throughout private and public spaces could lead to significant privacy concerns.

In the pre-IoT era, security cameras took the form of Closed Circuit Television Cameras (CCTVs). Research regarding the perceptions and behaviors surrounding CCTV can inform our understanding of the widespread use of cameras in smart spaces. CCTV surveillance cameras have been widely adopted by municipalities and businesses around the world to reduce crime and increase public safety. Studies suggest that CCTVs can lead to crime reduction in some cases, particularly for property crimes, and that camera surveillance is most suitable for small, well-defined areas [70], such as to reduce vehicle crimes in a parking garage.

Even when they are deployed in public spaces, CCTVs can raise a number of privacy concerns. One's autonomy and dignity can be reduced due to being under surveillance. Even when the presence of a CCTV camera is known, people typically cannot make a determination who is really behind that camera. (Not) knowing who is watching can influence how people behave. Surveillance can also have chilling effects on civil liberties and freedoms and can be particularly harmful to vulnerable populations, such as prisoners or students. Despite these concerns, the well-established use of CCTVs for public safety leads to different privacy perceptions and expectations compared to other camera-based technologies, such as smartphones or drones [71].

One challenge with CCTV is whether and how people are notified that they are under video surveillance. The most widely used way to inform people of CCTVs is to put up a sign indicating that people are within coverage of a camera. When they are clearly visible, even these notices themselves can increase the level of deterrence. However, in many cases such notices are far from effective since people rarely notice them or may become habituated to them over time. Surveillance notices also tend to provide little or no information about what happens with the captured recordings. Video technologies are also becoming smarter, with increasing capabilities toward facial and activity recognition. Again, though, surveillance notices tend to give little indication of the kind of processing that occurs, and there is typically no way for the public to access and control the data collected about them.

In recent years, IoT cameras have joined the ranks of CCTVs and are now being used throughout residential areas to provide for homeowners' security, but also collectively for neighborhood safety and security. While their motivation may be similar to CCTVs—to provide for the safety and security of one's home and belongings—this expansion of surveillance into more private spaces is likely to increase privacy risks. Privately owned IoT cameras are likely even less visible than CCTVs, with no notice at all to passersby. People will remain unaware of the extent to which they are being recorded as they drive down a road or walk down a sidewalk. Rather than prevent crime, knowledge of recording may have chilling effects on behavior in one's own private spaces. For example, residents may be less likely to speak freely in their own yard or to briefly step outside in a bathrobe if they expect to be recorded by a neighbor's camera. Finally, while cameras may be deployed by individuals on their own property, applications such as Citizens and Neighbors are

enabling the sharing of videos with neighbors and law enforcement [72, 73], thereby greatly expanding the potential audience for those videos.

In summary, the privacy challenges of security cameras are as follows:

- **Being recorded can change one's behavior** which can reduce crime but may also be perceived as a violation of one's privacy.
- **People are often unaware of, or get habituated to, surveillance notices.** Such notices also typically do not reveal the identity of the recipient or how they process the recordings, and they do not allow for access and control.
- **Privately owned IoT cameras further exacerbate these privacy issues,** as they tend to inconspicuously surveil more private spaces.

11.4 Solutions and Guidelines

Much of the research examining the privacy challenges and user perceptions in IoT have resulted in recommended design guidelines for supporting users' privacy needs through privacy features and interfaces. However, there has been considerably less research into how well different kinds of privacy controls could satisfy those guidelines or into novel mechanisms that specifically address the privacy challenges of IoT. In this section, we present these guidelines along with research into related solutions, including:

- Users need **greater awareness** of the data practices of IoT devices, which can be provided, in part, by additional **privacy notices**.
- Key **privacy controls** should be provided **on the device** itself to be easily accessible to all people in the environment.
- IoT devices and applications should implement state-of-the-art measures to maintain users' **data privacy**.
- IoT devices and applications should provide **flexible privacy controls** that give users adequate choices over the collection and sharing of their data.
- Users need **community-oriented privacy features** to support the many different kinds of users that interact within an IoT environment.
- **Context-adaptive privacy mechanisms** could reduce burden on users by personalizing settings and recommendations to the users and their context.

Providing adequate privacy choices is a challenge in IoT, due to the wide variety of devices, data, and contexts of use (resulting in a complex decision landscape) along with the absence of a dedicated user interface (resulting in limited opportunities for interaction). Thus, researchers and designers need to more fully examine the design space for providing various privacy mechanisms and controls. A good example is a recent paper by Feng et al., which introduced a design space for privacy choices in which they present five key dimensions of providing meaningful privacy controls in IoT [74]. These five dimensions include choice type, functionality, timing, channel, and modality. For example, in the timing dimension,

privacy choices can be delivered to users at six possible times: at setup, just in time, context-aware, periodic, on-demand, and personalized. In the remainder of the section, we provide a number of guidelines for supporting users' privacy needs, along with examples of privacy mechanisms attempting to address some of those guidelines. However, there is still significant need to expand upon these solutions to tackle the privacy issues raised above.

11.4.1 Privacy Notices and Awareness Mechanisms

One critical set of solutions to IoT privacy issues is to make users more aware of the privacy implications and risks of their interactions with IoT devices, so that they can make more informed privacy decisions. The guidelines addressing this need include the following:

- Provide privacy notices on the packaging and materials that come with the physical product, so that users can review data practices before they purchase the product and while they are first getting started setting up a device
- Provide privacy notices wherever the user may interact with the device, be that on the device, within an accompanying app, or in an online account
- Make privacy notices brief and focused around what the user would most care about or find surprising
- Make data collection and aggregation visible to the user as they interact with the device or accompanying application
- Provide periodic nudges regarding the data practices of the device, to allow users to learn more about data collection and reconsider those practices
- Provide mechanisms for users to discover what IoT devices are around them

Despite their limitations, the primary way that users learn details about the data practices of a device or application is through various privacy notices. One of the most common formats of privacy notice is the *privacy policy* available on most organizations' websites. However, current privacy policies are often lengthy and complex legal documents that contain detailed information related to a company's data practices. Research has identified many issues with such privacy policies, such as being hard to understand, time-consuming to read, and difficult to access [75, 76]. In the IoT context, many of these issues become even more prominent due to the nature of the physical devices. Unlike a website, where a privacy policy can be provided through a link on the web page, IoT devices generally have a very small screen, if they even have a screen at all. Instead, IoT device manufacturers require users to go to their product website to read privacy policies if desired, making it even more difficult for users to understand the data practices of IoT devices.

To combat this issue, Emami-Naeini et al. have proposed an IoT Security and Privacy Label [77]. The design of the label is inspired by the nutrition labels on

food packages, where the key nutrition information is conveyed to the consumer in a brief, standardized format. The IoT Security and Privacy Label is designed to be placed on the package of any IoT device and contains all the key information regarding the device's data practices (e.g., data collection purposes, data storage location, data sharing practices, etc.). This would allow users to read the label before purchase and compare the practices of similar products through the label. While these labels are not yet adopted by IoT manufacturers, privacy labels are beginning to be adopted in other domains. For example, Apple recently introduced a privacy label requirement for iOS apps, based on this and prior research [78].

Researchers have also examined how to provide privacy ratings or reviews to consumers to help them make purchase or use decisions. Consumers can find many different organizations, such as Consumer Reports, that review and rate all kinds of products on a variety of dimensions. In a similar vein, for privacy, Mozilla has created an online guide called "privacy not included" where consumers can learn about the data practices and possible risks from different smart home and IoT devices, so that existing users can assess their risk, and potential buyers can decide whether and which device to buy [79].

When privacy notices are salient and easily accessible to users, they can impact decision making. Yet, users are not likely to continue to view those notices as they interact with an IoT device. Thus, a critical solution is to make data collection and use visible within the interface of the device itself. This can be accomplished by various *data views* that show an aggregate of the collected information [80], along with detailed logs that can be accessed on demand. Yet, users may not always review such information, particularly if they do not regularly interact with the app that accompanies the IoT device. Thus, devices can also periodically nudge users regarding some aspect of their data collection, to prompt them to reflect on those data practices or review them in more detail. This has been investigated outside of IoT for mobile devices, for instance, where users were provided with periodic messages about how often different apps access their location [81].

Another major awareness challenge in IoT environments is how users can learn what active devices are nearby, particularly when they are in spaces that they do not control. Thus, another class of solutions helps bystanders discover IoT devices in their immediate surroundings. For example, *IoT Inspector* provides an easy way to understand what devices are connected in an IoT environment [82]. By scanning a user's network through a web app, IoT Inspector is able to identify all devices that are connected to the user's network and provide users with information such as device names, manufacturers, and IP addresses. For example, when a user stays at an AirBnB apartment, they learn of potential data collection within the apartment by scanning the network and identifying any connected IoT devices [83]. There are several other tools that provide somewhat similar functionalities, for example, IoT Sentinel [84] and Peek-a-Boo [85].

11.4.2 *On Device Controls*

Many users interact with smart devices through an accompanying mobile app. Yet, controls that are on the device itself are more accessible to everyone in the environment. Thus, guidelines advise to:

- Provide visible indicators on the device itself that indicate data collection is occurring
- Provide key controls on the device hardware that limit or turn off data collection

Page et al. found that people draw from different conceptual models when it comes to interacting with IoT devices: A person drawing from an *Agentic* perspective has a higher affinity to leverage non-haptic modes of interaction, whereas someone drawing from a *User-Centric* perspective prefers to use physical buttons. Manufacturers of consumer IoT devices have tried to cater to the latter group of users by providing a limited set of physical buttons on their smart devices, for example, Google Nest and Echo Dot each come with a physical button to disable their microphone [41]. Similarly, Facebook's Portal has camera covers and a camera disable button, in addition to the microphone disable button [42]. Hardware mechanisms have the benefit of being usable by anyone around the device, providing both control and an indication of the status of the device to bystanders. They also provide an added assurance of privacy—for example, users may trust that a physical cover over a camera truly prevents recording, rather than can a digital control indicating that the camera is off.

Researchers have also examined novel approaches that interfere with or impact a device's physical ability to collect data. For example, "Alias" is a separate add-on device that paralyzes a smart voice assistant by preventing it from listening and only activates the assistant with a custom wake word from the user [86]. Others have investigated the idea of obfuscating sound at the microphone as a privacy protection [87].

11.4.3 *Data Privacy*

Studies have found that IoT device users are generally concerned about the data these devices collect and desire additional measures to maintain their data privacy. The guidelines for supporting this need include the following:

- Make transmission and storage of data encrypted by default
- Store data anonymously when possible
- Where possible, provide users the option to process and store data locally (e.g., inside their device, app, or home network) instead of sending it to a remote server
- Make it more difficult for manufacturers/advertisers to make unwanted inferences through novel mechanisms such as adding noise to the data

Within IoT environments, data is constantly transmitted and stored by the device itself, by accompanying app, by the manufacturer in the cloud, and everywhere in between. This data traffic often contains sensitive information, which means that each storage and transmission point creates a risk of data leakage and data breaches. Thus, users expect that organizations are utilizing reasonable practices for protecting this data from attackers and third parties. A basic step is for data-centric encryption to be in place the moment the data is created within an IoT device and at every other point where it is stored. By enforcing encryption and making users aware of it, device manufacturers can gain users' trust as well as meet regulatory standards such as the California Consumer Privacy Act (CCPA).

Although encryption is a crucial step in ensuring the security and privacy of IoT data, recent research found that in-home activities such as sleep patterns, presence, and interaction with devices can be inferred even from encrypted IoT data using a technique called *traffic analysis* [21]. There have been several attempts to prevent such inferences from occurring. For instance, Hoof et al. secured a private messaging system from traffic analysis by shaping the traffic to a predetermined rate [88]. Apthorpe et al. introduce noise to shape the IoT network traffic to limit inference from traffic rate metadata [21]. Such mechanisms should be more extensively researched to encourage adoption.

Beyond encryption, research has found that many users desire to store and process their data locally rather than on a remote server [89, 90]. This solution gives more control to end users, who could explicitly choose to share data with manufacturers or other parties only when they want the benefit that this provides. Users could also choose to apply different levels of aggregation to their data before sharing it, thereby limiting the details of what is known or stored by others. For example, a fitness tracker user may not want to share their running route, but they may be willing to share how many miles they ran to receive some service (e.g., competing with friends). IoT users may also be more comfortable with data sharing when their data remains anonymous and cannot be linked back to their real identity. However, as outlined above, users often do not realize the extent to which inferences could occur, including those that could reidentify them from seemingly anonymous information. Introducing a carefully controlled amount of noise to the data can reduce the efficacy of such inferences, without significantly reducing the usefulness of the data for its intended purposes [91].

11.4.4 Community-Oriented Controls

IoT devices are regularly shared among a number of users. Therefore, designers must take a community-oriented view of IoT devices and applications, which includes providing features and controls that enable collective usage. Guidelines include to:

- Provide flexible and fine-grained sharing capabilities to allow users to share devices with many different kinds of people

- Make transparent what is accessible when devices or data are shared with other people
- Provide mechanisms to determine how different people are using devices or accessing data
- Learn the most prevalent sharing patterns and goals to support the design of sharing and access control capabilities

While many IoT applications allow users to share devices and data with others, many studies report that users find existing sharing capabilities too limited and request more fine-grained controls (e.g. [54]). Thus, a common guideline is to allow for more flexible and fine-grained sharing with different types of recipients. This is not only true for sharing data but also for control over the devices themselves. Another key limitation is that it can be challenging to determine exactly what is shared with whom, both at setup and over time. Thus, applications need mechanisms for users to be able to determine what other users will have access to, and be able to tell who is accessing those controls or that information over time.

Even within a household, different members may have different needs that are hard to monitor and control. For example, one particular study found that and 20% of kids aged 4–11 talk to their smart voice assistant for more than 5 h per week [92]. Children may get access to inappropriate content and reveal private information during their interactions. Parents already struggle with maintaining children's online safety with traditional devices; IoT devices make the situation even more difficult, as many reside in a common space designed to be used by all household members. To limit children's access to smart voice assistants, device manufacturers already provide parental control modes such as Amazon FreeTime and the Google Family App. Researchers have also examined ways to automatically determine content that is inappropriate for certain users within voice assistant conversations, such as Skillbot developed by Le et al. [93]. While parents may aim to protect their children from inappropriate content, designers must also protect the privacy of potentially vulnerable users from others within a smart space. For example, Freed et al. examined how technology can be exploited to enable intimate partner abuse [94]. Smart devices only provide increased capability for stalking and surveillance, which few have examined.

Users may desire a range of mechanisms to share access to their IoT devices depending on contextual factors. For example, in an attempt to improve the security and privacy tensions in a multi-user smart home, Zeng et al. developed an app that includes features such as location-based access controls, supervisory access controls, the ability to ask for permission (i.e., reactive access control), along with notifications on how others users are using a device [95]. However, in their field study, they found that users did not use many of the provided access control capabilities. Arguably, despite users' stated desire for more fine-grained access control features, supporting the full complexity of users' needs could result in interfaces that are more complex than some users are willing to utilize.

To combat this problem, designers must make an effort to understand users' most important goals and their most common interaction patterns regarding IoT

access control, so that they can support those goals and patterns more explicitly. For example, researchers have developed privacy-setting interfaces that are structured based on the relative importance of each contextual parameter [90, 96]. Likewise, Alqhatani et al. [23] described six particular sharing patterns of different audiences of fitness tracker information based on users' health and fitness goals. One of those patterns—sharing with healthcare providers—was not supported at all by existing devices.

11.4.5 Context-Adaptive and User-Tailored Privacy

One of the key challenges in privacy-preserving IoT is the contextual nature of privacy-related decisions and the explosion of contexts that are possible in this domain. Recent privacy regulations (e.g., California Consumer Privacy Act and General Data Protection Regulation, see Chap. 17) require by law that users can have more control over the data collected by IoT devices. Yet, providing controls that enable users to control the capture and sharing of information within every possible combination of context is simply too overwhelming. A potential solution to this problem is to provide context-adaptive and user-tailored privacy controls. More details about this solution can be found in Chap. 16; here we provide IoT-specific guidelines:

- Study the context-dependency of users' IoT privacy decisions
- Use machine learning to predict users' privacy preferences, for example, by creating comprehensive privacy profiles
- Automate the privacy practices of IoT devices based on the context and the profile of the user

A large number of studies have demonstrated that users' comfort with IoT privacy practices vary across different factors such as the type of data recorded, the location where it is recorded, who the data is shared with, the perceived value of the data, and the benefits provided by services using that data [5, 6, 63, 90, 96–101]. In the context of public IoT, Naeini et al. [5] used vignettes to study many of these factors with over 380 different use cases across 1000 users. Their results indicate that people are most uncomfortable when data is collected in their home and prefer to be notified when such collection occurs. Similarly, a survey study by Lee and Kobsa [6] found that monitoring of users' personal spaces, such as their homes, was not acceptable to participants, as well as monitoring performed by the government or unknown entities. Other studies have found that people are most concerned with certain types of data, namely videos, photos, and bio-metric information, particularly when this information is gathered inside the home [5, 6, 14, 102, 103].

In a smart home setting, He et al. [90] find that when IoT devices share data with one another, users are most concerned about where that data is stored, followed by the types of devices that act as the sender and recipient of the information and the purpose of the data transmission. Significant interaction effects between sender,

recipient, and purpose suggest that users have complex preferences that depend on multiple contextual parameters at once. Likewise, Barbosa et al. concluded that people's privacy perceptions regarding smart home IoT devices depend on not only the data types but also the purpose of data collection and who collect the data [104]. In another large vignette study, Apthorpe et al. [99] found that participants' acceptance of data collection and sharing was dependent on both the recipient of the information and the specific conditions under which the information was shared. Their results also suggest that users' privacy norms may change with the continued use of specific devices. Results of a different vignette survey by Horne et al. [105] suggest that those changes are not always toward more acceptance of data-sharing.

Beyond context-dependency, IoT privacy decisions also vary significantly by user [90, 96], suggesting that users must be able to make a personal decision as to whether a certain scenario warrants the collection and/or sharing of information. Unfortunately, the sheer number of contextual parameters to consider in this decision will likely substantially increase the complexity of the privacy-setting interfaces of IoT devices. In response, researchers have attempted to reduce the apparent variety of IoT privacy decisions to a small set of concise *privacy profiles* for users to choose from [90, 96, 106], thereby reducing the complexity of the privacy-setting task.

Taking this approach one step further, researchers have proposed frameworks to support IoT privacy decision-making by adapting the privacy settings to varying privacy contexts and/or by recommending users to make certain privacy decisions [6, 106–108]. More details about this user-tailored approach can be found in Chap. 16.

11.5 Conclusion

This chapter has covered the prevailing privacy challenges of IoT environments from a user-centered perspective. We have demonstrated that the introduction of sensor-based Internet-connected technologies in real-world environments—be they wearables, household devices, or devices in the public domain—exacerbate existing issues with online privacy and physical privacy and introduce new, unique challenges as well.

These challenges exist because IoT devices can inconspicuously collect vast amounts of data and perform inferences on this data to paint a detailed picture of the preferences and activities of their users (and even of bystanders). The typical lack of a comprehensive user interface reduces users' awareness of these data practices and limits their control over them—if control is even provided at all.

The resulting privacy issues are further exacerbated by the fact that existing IoT devices tend to offer limited configuration of how devices and data are shared among a community of users as well as inadequate security protections to prevent outsiders from gaining unwanted access to sensitive data and/or functionality of the device.

The final section of our chapter offers solutions and guidelines for IoT researchers and developers to increase users' awareness about and control over their privacy, both regarding the data practices of the IoT manufacturer and the use of IoT devices by multiple users. While we advocate for granular, on-device control, we also acknowledge that fine-grained control can be daunting to users. Context-adaptive and user-tailored privacy solutions may provide IoT users with adequate control over their privacy while at the same time reducing the burden of effecting this control.

The unrelenting evolution of artificial intelligence and sensor technologies, paired with the continuing miniaturization of processing and networking chips, suggests that current IoT technologies only scratch the surface of what future technologies in this realm will be capable of. For privacy researchers, it is therefore important to "future-proof" their research by not just focusing on what is currently possible with IoT technologies but to anticipate the socio-technical consequences of the imaginable. Likewise, for developers and manufacturers, it is important to acknowledge that many of the IoT devices of today will operate alongside the ones that will be created in the future, and to design the privacy mechanisms of today's systems accordingly. We hope that the tremendous benefits promised by the rise of IoT will be paired with a powerful user experience that respects the privacy of users and bystanders alike.

References

1. Ashton, K., et al. 2009. That 'internet of things' thing. *RFID Journal* 22 (7): 97–114.
2. Lee, I., and K. Lee. 2015. The internet of things (IoT): Applications, investments, and challenges for enterprises. *Business Horizons* 58 (4): 431–440.
3. Haghi, M., K. Thurow, and R. Stoll. 2017. Wearable devices in medical internet of things: scientific research and commercially available devices. *Healthcare Informatics Research* 23 (1): 4.
4. Motti, V.G., and K. Caine. 2015. Users' privacy concerns about wearables. In *Financial Cryptography and Data Security*, 231–244. Berlin: Springer.
5. Naeini, P.E., S. Bhagavatula, H. Habib, M. Degeling, L. Bauer, L.F. Cranor, and N. Sadeh. 2017. Privacy expectations and preferences in an iot world. In *Thirteenth Symposium on Usable Privacy and Security (SOUPS) 2017*, 399–412.
6. Lee, H. and A. Kobsa. 2016. Understanding user privacy in internet of things environments. In *2016 IEEE 3rd World Forum on Internet of Things (WF-IoT)*, 407–412. Piscataway: IEEE.
7. Nyc midtown congestion management system. https://www1.nyc.gov/html/dot/html/pr2012/pr12_25.shtml. Accessed 09 Nov 2020.
8. STAFF, W. 2020. Georgia city moves forward with extensive water loss control program. <https://waterfm.com/georgia-city-moves-forward-with-extensive-water-loss-control-program/>. Accessed 09 Nov 2020.
9. Bloom, C., J. Tan, J. Ramjohn, L. Bauer. 2017. Self-driving cars and data collection: Privacy perceptions of networked autonomous vehicles. In *Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)*, Santa Clara, CA, 357–375. San Francisco Bay: USENIX Association.

10. Law, B.H. 2021 What you need to know about driverless cars and privacy. <https://medium.com/@baumhedlund/what-you-need-to-know-about-driverless-cars-and-privacy-8720d46e8877>. Accessed 04 Nov 2021.
11. Tabassum, M., J. Kropczynski, P. Wisniewski, and H.R. Lipford. 2020. Smart home beyond the home: A case for community-based access control. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI '20, 1–12. New York: Association for Computing Machinery.
12. Zheng, S., N. Apthorpe, M. Chetty, and N. Feamster. 2018. User perceptions of smart home iot privacy. *Proceedings of the ACM on Human-Computer Interaction 2* (CSCW): 200:1–200:20. <https://doi.org/10.1145/3274469>. <http://doi.acm.org/10.1145/3274469>
13. Vitak, J., Y. Liao, P. Kumar, M. Zimmer, and K. Kritikos. 2018. Privacy attitudes and data valuation among fitness tracker users. In *iConference*.
14. Lee, L., J. Lee, S. Egelman, and D. Wagner. 2016. Information disclosure concerns in the age of wearable computing. In *NDSS Workshop on Usable Security (USEC)*, vol. 1.
15. Peppet, R. 2014. Regulating the internet of things: First steps toward managing discrimination, privacy, security, and consent. *Texas Law Review* 93: 85–179.
16. Emami-Naeini, P., H. Dixon, Y. Agarwal, and L.F. Cranor. 2019. Exploring how privacy and security factor into iot device purchase behavior. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, New York, NY, 1–12. New York: Association for Computing Machinery.
17. Zeng, E., S. Mare, and F. Roesner. 2017. End user security and privacy concerns with smart homes. In *Thirteenth Symposium on Usable Privacy and Security (SOUPS) 2017*, 65–80.
18. Kröger, J. 2018. Unexpected inferences from sensor data: a hidden privacy threat in the internet of things. In *IFIP International Internet of Things Conference*, 147–159. Berlin: Springer.
19. Wang, H., T.T.-T. Lai, and R. Roy Choudhury. 2015. Mole: Motion leaks through smartwatch sensors. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*, MobiCom '15, 155–166. New York: Association for Computing Machinery.
20. Srinivasan, V., J. Stankovic, and K. Whitehouse. 2008. Protecting your daily in-home activity information from a wireless snooping attack. In *Proceedings of the 10th International Conference on Ubiquitous Computing*, UbiComp '08, 202–211. New York: Association for Computing Machinery.
21. Apthorpe, N., H.D. Yuxing, R. Dillon, N. Arvind and F. Nick. 2019. Keeping the smart home private with smart(er) IoT traffic shaping. *Proceedings on Privacy Enhancing Technologies*, 2019 (3): 128–148. <https://doi.org/10.2478/popets-2019-0040>
22. House votes to allow internet service providers to sell, share your personal information. <https://www.consumerreports.org/consumerist/house-votes-to-allow-internet-service-providers-to-sell-share-your-personal-information/>. Accessed 26 Nov 2019.
23. Alqhatani, A., and H.R. Lipford. 2019. “there is nothing that i need to keep secret”: Sharing practices and concerns of wearable fitness data. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*. Santa Clara: USENIX Association.
24. Gerber, N., B. Reinheimer, and M. Volkamer. 2018. Home sweet home? investigating users' awareness of smart home privacy threats. In *Proceedings of an Interactive Workshop on the Human Aspects of Smarthome Security and Privacy (WSSP)*, Baltimore, MD, August 12, 2018. USENIX.
25. Home assistant adopter beware: Google, amazon digital assistant patents reveal plans for mass snooping. <https://www.consumerwatchdog.org/privacy-technology/home-assistant-adopter-beware-google-amazon-digital-assistant-patents-reveal>. Accessed 26 Nov 2019.
26. McStay, A. 2016. Empathic media and advertising: Industry, policy, legal and citizen perspectives (the case for intimacy). *Big Data & Society* 3 (2): 2053951716666868.
27. The facebook and cambridge analytica scandal, explained with a simple diagram. <https://www.vox.com/policy-and-politics/2018/3/23/17151916/facebook-cambridge-analytica-trump-diagram>. Accessed 26 Nov 2019.

28. Cambridge analytica: how did it turn clicks into votes? <https://www.theguardian.com/news/2018/may/06/cambridge-analytica-how-turn-clicks-into-votes-christopher-wylie>. Accessed 13 April 2021.
29. Trust in facebook has dropped by 66 percent since the cambridge analytica scandal. <https://www.nbcnews.com/business/consumer/trust-facebook-has-dropped-51-percent-cambridge-analytica-scandal-n867011>. Accessed 13 April 2021.
30. Smart streetlights program. <https://www.sandiego.gov/sustainability/energy-and-water-efficiency/programs-projects/smart-city>. Accessed 09 Nov 2020.
31. Nyc connected vehicle project. <https://cvt.nyc/>. Accessed 09 Nov 2020.
32. Garg, R., and C. Moreno. 2019. Understanding motivators, constraints, and practices of sharing internet of things. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 3 (2): 1–21.
33. He, W., M. Golla, R. Padhi, J. Ofek, M. Dürmuth, E. Fernandes, and B. Ur. 2018. Rethinking access control and authentication for the home internet of things (IoT). In *27th USENIX Security Symposium (USENIX Security 18)*, Baltimore, MD, 255–272. Berkeley: USENIX Association.
34. Geeng, C., and F. Roesner. 2019. Who’s in control? interactions in multi-user smart homes. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI ’19, 1–13. New York: Association for Computing Machinery.
35. Page, X., P. Bahirat, M.I. Safi, B.P. Knijnenburg, and P. Wisniewski. 2018. The internet of what? understanding differences in perceptions and adoption for the internet of things. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2 (4): 1–22.
36. Desai, B.C. 2017. Iot: imminent ownership threat. In *Proceedings of the 21st International Database Engineering & Applications Symposium*, 82–89.
37. Janeček, V. 2018. Ownership of personal data in the internet of things. *Computer Law & Security Review* 34 (5), 1039–1052.
38. Mare, S., F. Roesner, and T. Kohno. 2020. Smart devices in airbnbs: Considering privacy and security for both guests and hosts. *Proceedings on Privacy Enhancing Technologies* 2020 (2): 436–458.
39. Google glass users fight privacy fears. <https://www.cnn.com/2013/12/10/tech/mobile/negative-google-glass-reactions>. Accessed 13 April 2021.
40. Lau, J., B. Zimmerman, and F. Schaub. 2018. Alexa, are you listening?: Privacy perceptions, concerns and privacy-seeking behaviors with smart speakers. *Proceedings of the ACM on Human-Computer Interaction* 2 (CSCW): 102:1–102:31.
41. Johnson, D. 2019. How to stop your Google Home from listening to you and storing your audio data. <https://www.businessinsider.com/how-to-stop-google-home-from-listening-to-me>
42. PortalPrivacy, <https://portal.facebook.com/privacy>
43. Bertino, E., and N. Islam. 2017. Botnets and internet of things security. *Computer* 50 (2): 76–79.
44. Copos, B., K. Levitt, M. Bishop, and J. Rowe. 2016. Is anybody home? inferring activity from smart home network traffic. In *2016 IEEE Security and Privacy Workshops (SPW)*, 245–251. Piscataway: IEEE.
45. Leyden, J. 2015. Samsung smart fridge leaves Gmail logins open to attack. https://www.theregister.com/2015/08/24/smart_fridge_security_fubar/ Accessed October 4, 2021.
46. Goodin, D. 2015. Baby monitors wide open to hacks that expose users’ most private moments. *ars technica*. 2015. <https://arstechnica.com/information-technology/2015/09/9-baby-monitors-wide-open-to-hacks-that-expose-users-most-private-moments/>
47. Dickson, B. 2015. Why IoT security is so critical. <https://trn.ch/314FL4Z>. Accessed 4 October 2021
48. Mantas, G., Lymberopoulos, D., Komninos, N. 2010. Security in smart home environment. In *Wireless Technologies for Ambient Assisting Living and Healthcare: Systems and Applications; IGI Global: Hershey, PA, USA, pp. 170–191*. DOI: 10.4018/978-1-61520-805-0.ch010

49. Braun, L. 2016. Human centered security: (How) can the typical smart home user make his home more secure?, *ResearchGate*, Aug. 2016. https://www.researchgate.net/publication/305850389_Human_Centered_Security_How_can_the_typical_smart_home_user_make_his_home_more_secure
50. Iot traffic in the enterprise is rising. so are the threats. <https://www.zscaler.com/blogs/security-research/iot-traffic-enterprise-rising-so-are-threats>. Accessed 13 April 2021.
51. Most iot devices are an attack waiting to happen, unless manufacturers update their kernels. <https://www.techrepublic.com/article/most-iot-devices-are-an-attack-waiting-to-happen-unless-manufacturers-update-their-kernels/>. Accessed 13 April 2021.
52. Lowens, B., V.G. Motti, and K. Caine. 2017. Wearable privacy: Skeletons in the data closet. In *2017 IEEE International Conference on Healthcare Informatics (ICHI)*, 295–304. <https://doi.org/10.1109/ICHI.2017.29>
53. Dong, M., L. Chen, and L. Wang. 2019. Investigating the user behaviors of sharing health- and fitness-related information generated by mi band on weibo. *International Journal of Human—Computer Interaction* 35 (9): 773–786.
54. Lowens, B.M. 2018. Toward privacy enhanced solutions for granular control over health data collected by wearable devices. In *Proceedings of the 2018 Workshop on MobiSys 2018 Ph.D. Forum*, MobiSys PhD Forum '18, 5–6. New York: Association for Computing Machinery. <https://doi.org/10.1145/3212711.3212714>
55. From cheating to pregnancy reveals, wearables know what you're doing intimately. <https://www.inverse.com/mind-body/from-cheating-to-pregnancy-reveals-wearables-know-what-you-are-doing-intimately>. Accessed 13 April 2021.
56. Meteriz, Ü., Fazıl Yıldırım, N., Kim J, and D. Mohaisen. 2020. Understanding the potential risks of sharing elevation information on fitness applications. *2020 IEEE 40th International Conference on Distributed Computing Systems (ICDCS)*, 464–473, doi: [10.1109/ICDCS47774.2020.00063](https://doi.org/10.1109/ICDCS47774.2020.00063).
57. Fitness tracking app strava gives away location of secret us army bases. <https://www.theguardian.com/world/2018/jan/28/fitness-tracking-app-gives-away-location-of-secret-us-army-bases>. Accessed 13 April 2021.
58. Rader, E., and J. Slaker. 2017. The importance of visibility for folk theories of sensor data. In *Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)*, 257–270. Santa Clara: USENIX Association.
59. Gabriele, S., and S. Chiasson (2020). Understanding fitness tracker users' security and privacy knowledge, attitudes and behaviours. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI '20, 1–12. New York: Association for Computing Machinery.
60. 'Alexa, are you invading my privacy?' The dark side of our voice assistants. (2019). <http://www.theguardian.com/technology/2019/oct/09/alexa-are-you-invading-my-privacy-the-dark-side-of-our-voice-assistants> Section: Technology.
61. Amazon customer receives 1,700 audio files of a stranger who used alexa. <https://www.npr.org/2018/12/20/678631013/amazon-customer-receives-1-700-audio-files-of-a-stranger-who-used-alexa?t=1570014709519&t=1570530199090>. Accessed 13 April 2021.
62. Cowan, B.R., N. Pantidi, D. Coyle, K. Morrissey, P. Clarke, S. Al-Shehri, D. Earley, and N. Bandeira. 2017. "what can i help you with?": Infrequent users' experiences of intelligent personal assistants. In *Proceedings of the 19th International Conference on Human-Computer Interaction with Mobile Devices and Services*, MobileHCI '17, New York: Association for Computing Machinery.
63. Malkin, N., J. Bernd, M. Johnson, and S. Egelman. 2018. "what can't data be used for?" Privacy expectations about smart tvs in the us. In *European Workshop on Usable Security (Euro USEC)*.
64. Hey, i didn't order this dollhouse! 6 hilarious alexa mishaps. <https://www.digitaltrends.com/home/funny-accidental-amazon-alexa-ordering-stories/>. Accessed 13 April 2021.
65. Nationwide Insurance to Give a Million Customers Echo Auto, Doubling Amazon's In-Car User Base. (2019). <https://voicebot.ai/2019/10/09/nationwide-insurance-to-give-a-million-customers-echo-auto-doubling-amazons-in-car-user-base/>. Section: Alexa skills.

66. Home assistant adopter beware: Google, amazon digital assistant patents reveal plans for mass snooping. <https://www.consumerwatchdog.org/privacy-technology/home-assistant-adopter-beware-google-amazon-digital-assistant-patents-reveal>. Accessed 13 April 2021.
67. Tabassum, M., T. Kosiński, A. Frik, N. Malkin, P. Wijesekera, S. Egelman, and H.R. Lipford. 2019. Investigating users' preferences and expectations for always-listening voice assistants. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 3 (4): 1–23.
68. Dear, K., K. Dutton, and E. Fox. 2019. Do 'watching eyes' influence antisocial behavior? a systematic review & meta-analysis. *Evolution and Human Behavior* 40 (3): 269–280.
69. Mazerolle, L., D. Hurley, and M. Chamlin. 2002. Social behavior in public space: An analysis of behavioral adaptations to cctv. *Security Journal* 15 (3): 59–75.
70. McLean, S.J., R.E. Worden, and M. Kim. 2013. Here's looking at you: An evaluation of public cctv cameras and their effects on crime and disorder. *Criminal Justice Review* 38 (3): 303–334. <https://doi.org/10.1177/0734016813492415>
71. Wang, Y., H. Xia, Y. Yao, and Y. Huang. 2016. Flying eyes and hidden controllers: A qualitative study of people's privacy perceptions of civilian drones in the us. *Proceedings on Privacy Enhancing Technologies* 2016 (3): 172–190.
72. Brush, A., J. Jung, R. Mahajan, and F. Martinez. 2013. Digital neighborhood watch: Investigating the sharing of camera data amongst neighbors. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work*, 693–700. Ne York: ACM.
73. Ring. Ring neighborhood watch. <https://shop.ring.com/pages/neighbors>. Accessed 09 Nov 2020.
74. Feng, Y., Y. Yao, and N. Sadeh. 2021. A design space for privacy choices: Towards meaningful privacy control in the internet of things. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21, New York, NY, 1–16. New York: Association for Computing Machinery. <https://doi.org/10.1145/3411764.3445148>
75. Luger, E., S. Moran, and T. Rodden. 2013. Consent for all: Revealing the hidden complexity of terms and conditions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, 2687–2696. New York: Association for Computing Machinery.
76. McDonald, A.M. and L.F. Cranor. 2008. The cost of reading privacy policies. *Isjlp* 4: 543.
77. Emami-Naeini, P., Y. Agarwal, L.F. Cranor, and H. Hibshi. 2020. Ask the experts: What should be on an iot privacy and security label? In *2020 IEEE Symposium on Security and Privacy (SP)*, 447–464. Piscatawy: IEEE.
78. What we learned from apple's new privacy labels. <https://nytimes.com/2021/01/27/technology/personaltech/apple-privacy-labels.html>nytimes.com/2021/01/27/technology/personaltech/apple-privacy-labels.html. Accessed 13 April 2021.
79. Mozilla - *privacy not included. <https://foundation.mozilla.org/en/privacynotincluded/>. Accessed 13 April 2021.
80. Wilkinson, D., P. Bahirat, M. Namara, J. Lyu, A. Alsubhi, P. Wisniewski, and B. Knijnenburg. 2019. Privacy at a glance: Exploring the effectiveness of screensavers to improve privacy awareness. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*. Under Review. New York: ACM.
81. Almuhiemedi, H., F. Schaub, N. Sadeh, I. Adjerid, A. Acquisti, J. Gluck, L.F. Cranor, and Y. Agarwal. 2015. Your location has been shared 5,398 times! a field study on mobile app privacy nudging. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 787–796.
82. Huang, D.Y., N. Apthorpe, F. Li, G. Acar, and N. Feamster. 2020. Iot inspector: Crowdsourcing labeled network traffic from smart home devices at scale. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4 (2): 1–21.
83. Yao, Y., J.R. Basdeo, O. R. McDonough, and Y. Wang. 2019. Privacy perceptions and designs of bystanders in smart homes. *Proceedings of the ACM on Human-Computer Interaction* 3 (CSCW): 1–24.

84. Miettinen, M., S. Marchal, I. Hafeez, N. Asokan, A.-R. Sadeghi, and S. Tarkoma. 2017. Iot sentinel: Automated device-type identification for security enforcement in IoT. In *2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS)*, 2177–2184. Piscataway: IEEE.
85. Acar, A., H. Fereidooni, T. Abera, A.K. Sikder, M. Miettinen, H. Aksu, M. Conti, A.-R. Sadeghi, and S. Uluagac. 2020. Peek-a-boo: I see your smart home activities, even encrypted! In *Proceedings of the 13th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, 207–218.
86. Project alias. http://bjoernkarmann.dk/project_alias. Accessed 08 Jan 2021.
87. Chandrasekaran, V., T. Linden, K. Fawaz, B. Mutlu, and S. Banerjee. 2018. Blackout and obfuscator: An exploration of the design space for privacy-preserving interventions for voice assistants. Preprint arXiv:1812.00263.
88. van den Hooff, J., D. Lazar, M. Zaharia, and N. Zeldovich. 2015. Vuvuzela: Scalable private messaging resistant to traffic analysis. In *Proceedings of the 25th Symposium on Operating Systems Principles, SOSP '15*, 137–152. New York: Association for Computing Machinery.
89. Yao, Y., J.R. Basdeo, S. Kaushik, and Y. Wang. 2019. Defending my castle: A co-design study of privacy mechanisms for smart homes. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, CHI '19*, 1–12. New York: Association for Computing Machinery.
90. He, Y., P. Bahirat, B.P. Knijnenburg, and A. Menon. 2019. A data-driven approach to designing for privacy in household iot. *ACM Transactions on Interactive Intelligent Systems (TiIS)* 10 (1): 1–47.
91. Chow, R., H. Jin, B. Knijnenburg, and G. Saldamli. 2013. Differential data analysis for recommender systems. In *Proceedings of the 7th ACM Conference on Recommender Systems, RecSys '13*, 323–326. New York: Association for Computing Machinery. <https://doi.org/10.1145/2507157.2507190>
92. Kids are spending more time with voice, but brands shouldn't rush to engage them. <https://www.emarketer.com/content/kids-are-spending-more-time-with-voice-but-brands-shouldnt-rush-to-engage-them>. Accessed 13 April 2021.
93. Le, T., D. Huang, N.J. Apthorpe, and Y. Tian. 2021. Skillbot: Identifying risky content for children in alexa skills. ArXiv abs/2102.03382.
94. Freed, D., J. Palmer, D.E. Minchala, K. Levy, T. Ristenpart, and N. Dell. 2017. Digital technologies and intimate partner violence: A qualitative analysis with multiple stakeholders. *Proceedings of the ACM on Human-Computer Interaction* 1 (CSCW): 1–22.
95. Zeng, E., and F. Roesner. 2019. Understanding and improving security and privacy in multi-user smart homes: a design exploration and in-home user study. In *28th {USENIX} Security Symposium ({USENIX} Security 19)*, 159–176.
96. Bahirat, P., Y. He, A. Menon, and B. Knijnenburg. 2018. A data-driven approach to developing iot privacy-setting interfaces. In *23rd International Conference on Intelligent User Interfaces*, 165–176.
97. Klasnja, P., S. Consolvo, J. Jung, B.M. Greenstein, L. LeGrand, P. Powledge, and D. Wetherall (2009). "when i am on wi-fi, i am fearless" privacy concerns & practices in everyday wi-fi use. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1993–2002.
98. Ghiglieri, M., M. Volkamer, and K. Renaud. 2017. Exploring consumers' attitudes of smart tv related privacy risks. In *International Conference on Human Aspects of Information Security, Privacy, and Trust*, 656–674. Berlin: Springer.
99. Apthorpe, N., Y. Shvartzshnaider, A. Mathur, D. Reisman, and N. Feamster. 2018. Discovering smart home internet of things privacy norms using contextual integrity. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2 (2): 1–23.
100. Lederer, S., J. Mankoff, A.K. Dey. 2003. Who wants to know what when? privacy preference determinants in ubiquitous computing. In *CHI'03 Extended Abstracts on Human Factors in Computing Systems*, 724–725.

101. Choe, E.K., S. Consolvo, J. Jung, B. Harrison, and J.A. Kientz. 2011. Living in a glass house: a survey of private moments in the home. In *Proceedings of the 13th International Conference on Ubiquitous Computing*, 41–44.
102. Aleisa, N., and K. Renaud. 2017. Yes, i know this iot device might invade my privacy, but i love it anyway! a study of saudi arabian perceptions. In *IoT BDS 2017: 2nd International Conference on Internet of Things: Big Data and Security, Porto*.
103. Das, A., M. Degeling, X. Wang, J. Wang, N. Sadeh, and M. Satyanarayanan. 2017. Assisting users in a world full of cameras: A privacy-aware infrastructure for computer vision applications. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 1387–1396. Piscataway: IEEE.
104. Barbosa, N.M., J.S. Park, Y. Yao, and Y. Wang. 2019. " what if?" predicting individual users' smart home privacy preferences and their changes. *PoPETs 2019* (4): 211–231.
105. Horne, C., B. Darras, E. Bean, A. Srivastava, and S. Frickel. 2015. Privacy, technology, and norms: The case of smart meters. *Social Science Research* 51: 64–76.
106. Sanchez, O.R., I. Torre, Y. He, and B.P. Knijnenburg. 2019. A recommendation approach for user privacy preferences in the fitness domain. *User Modeling and User-Adapted Interaction* 30: 513–565. <https://doi.org/10.1007/s11257-019-09246-3>
107. Schaub, F., B. Könings, M. Weber, and F. Kargl. 2012. Towards context adaptive privacy decisions in ubiquitous computing. In *2012 IEEE International Conference on Pervasive Computing and Communications Workshops*, 407–410. Piscataway: IEEE.
108. Sanchez, O.R., I. Torre, and B.P. Knijnenburg. 2020. Semantic-based privacy settings negotiation and management. *Future Generation Computer Systems* 111: 879–898. <https://doi.org/https://doi.org/10.1016/j.future.2019.10.024>, <https://www.sciencedirect.com/science/article/pii/S0167739X18317035>

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/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

