



Department for
Digital, Culture,
Media & Sport

Systematic Review of Security and Privacy Recommendations for non- mobile apps and app stores

October 2022



This systematic review was carried out by:



Dr. Mariam Elgabry

Bronic Ltd.

Professor Shane Johnson

University College London

Executive Summary

The aim of this review is to further inform the Department for Digital, Culture, Media and Sport's work on app security and privacy by conducting a review of recommendations suggested to improve the security and privacy of non-mobile apps and non-mobile app stores. This includes those linked to devices such as smart TVs, fitness, gaming and voice assistants. We used a systematic search methodology to review the academic literature, and supplemented this with a review of open-source materials and "grey" literature that is not disseminated by academic publishers.

A comprehensive search of articles published in or after 2017 until the 1st of July 2022 found that there have been exceedingly few studies (ten) that focus on improvements to non-mobile apps and app stores. The search was conducted using search engines that index over 60 databases – databases that specifically cover the information security, and computing literatures, as well as the wider set of academic disciplines. Across these ten articles, authors identified a total of 11 types of security threats for which a total of ten recommendations were suggested. These recommendations were provided for app stores, app developers, device manufacturers, and users.¹ However, only two of the identified recommendations were explicitly for app stores and only three were explicitly aimed at app developers, highlighting a significant gap in the literature that should be addressed.

¹ Some of them were sub-recommendations, and some were not about apps or app stores in particular, but were relevant to the review.

1. Introduction

A recent literature review² of security and privacy policies for mobile phone applications (hereafter, apps) and app stores, conducted for the Department for Digital, Culture, Media & Sport (DCMS), reported that some apps and mobile app stores offer a permissive environment which leads to malicious or risky apps being available to users. Additionally, the review's author found that there are notable variations in how different app stores guide and support app developers, maintain security and resolve issues. In light of these findings, and to help protect consumers from online threats, the Government is proposing a voluntary Code of Practice for all app stores and developers.

The aim of this review is to complement the above work by considering apps and app stores intended for non-mobile devices, such as gaming consoles and smart televisions. To address this issue, we conducted a systematic literature review to identify recommendations for improving the security of apps available through non-mobile app stores, and to identify issues that may be of relevance to DCMS's work on cyber security in these contexts. We used a systematic search methodology to review the academic literature (e.g., Petticrew & Roberts, 2006), and supplemented this with a review of other open-source materials and "grey" literature that is not disseminated by academic publishers.

The key question addressed by this review was: *what recommendations have been suggested in the literature to address security and privacy issues that are linked to and might be facilitated by non-mobile apps and app stores?*

The review covers recommendations applicable to app users, and where available, app developers and device vendors.

1.1 Structure of this report

This report is organised as follows. Section 1.2 defines what is considered a non-mobile app/app store in this report and outlines the range of different devices currently in the market. For transparency, subsection 1.3 details the review methodology employed. Section 2 discusses the recommendations identified in the literature and briefly summarises the types of threats these were intended to address. The report concludes with a summary of the recommendations.

1.2. Background

As discussed, this report specifically considers non-mobile apps and app stores. It therefore excludes smartphone and tablet devices, and platforms such as Android, iOS and the *Windows Store* (which were considered in Furnell, 2021). Like mobile apps, non-mobile apps that are used by smart devices (e.g. Smart TV's) also have app stores (e.g., the SmartThings Marketplace³). These stores offer their own apps and also accept apps from third-party developers' which may, for example, enable home-automation using cloud frameworks and the integration of multiple home IoT devices. The latter can range from sensors (e.g. a microphone, or a movement detector) to large digital appliances and can enable complicated "If This Then That" operations across devices (e.g., "If an outdoor sensor is triggered, turn the outside lights on and turn on an external video camera") to be performed by a set of apps and devices. Importantly, Smart applications can control physical devices that are implicated

² Furnell (2021). Literature review on security and privacy policies in apps and app stores. DCMS. <https://www.gov.uk/government/consultations/app-security-and-privacy-interventions/literature-review-on-security-and-privacy-policies-in-apps-and-app-stores>

³ SMARTTHINGS, INC. SmartThings Marketplace. <https://support.smartthings.com/hc/en-us/articles/205379924-Marketplace>, 2016.

in security (e.g., door locks) and/or health (e.g., through fitness wearables) that could cause serious harm if they are not appropriately protected and hence it is important that such applications meet (at least) adequate security standards. In the following subsection, we discuss the range of different devices currently in the market that were covered by our search.

1.2.1 Types of devices, apps and app stores

Everyday activities have been redefined with the introduction of internet connected devices, often referred to as the Internet of Things (IoT). These include devices for the “smart home” such as smart TV’s, gaming consoles, and voice assistants (e.g., Amazon Alexa). They also include devices such as smartwatches and headset wearables that have been developed for lifestyle purposes such as fitness and meditation.

Consumers can easily download and install apps for such devices from vendor-specific app markets, or even develop their own. For any type of device, there are differences in the types of app or app stores used. Official apps and app stores are those provided directly by the manufacturer of the device. In contrast, third-party apps and app stores are not. On the one hand, the latter introduces accessibility and freedom for the developer and end user, and can lead to the development of a much wider array of apps than might otherwise emerge, encouraging innovation. On the other hand, as these apps and the stores through which they are provided are not controlled by a single organisation, this may reduce the oversight and governance of what is being offered, which could lead to differences in the attention paid to privacy and security and the effectiveness of any approaches taken to preserve them. Such variation implies the need for clear recommendations (and codes of practice) that can be followed across the entire ecosystem. The types of non-mobile apps and app stores considered here are those associated with smart consumer devices and they include both official apps/app stores for smart home systems such as Samsung SmartThings, Amazon Alexa, and LIFX SmartLight, as well as third-party apps/app stores.

1.3 Review Methodology

1.3.1 Databases and Search Terms Used

The final search was conducted on 1 July 2022, and to identify relevant articles we searched the academic electronic databases ProQuest Central⁴, ACM digital library⁵ and IEEE Xplore⁶. Collectively, these provide comprehensive coverage of published academic research across the social, engineering and physical sciences, as well as the computing and information security literatures. General web searches were also conducted to identify relevant reports and media coverage of known incidents of interest. ProQuest Central indexes media reports as well as the academic literature, but we also conducted a general web search using Google Search to provide more extensive coverage. For

⁴ ProQuest Central is a comprehensive search engine for academic literature. It covers 60 Databases across all major subject areas, including business, health and medical, social sciences, science, and technology. It indexes full-text scholarly journals, Newspapers, magazines, Dissertations, working papers, case studies, and Market reports. The portfolio of databases covered includes the Criminal Justice Database, Computing Database, Library Science Database, Science Database, Social Science Database, Psychology Database and continent- specific databases covering technology and social sciences (such as the Australia & New Zealand Database, Continental Europe Database, East & South Asia Database, East Europe & Central Europe database etc.), and ProQuest Dissertations & Theses Global. It also includes the Association for Computing Machinery (ACM) digital library.

⁵ ACM digital library is a comprehensive database of full-text articles and bibliographic literature covering computing and information technology from the Association for Computing Machinery publications.

⁶ IEEE Xplore Digital Library is an indexed database of articles and papers on computer science, electrical engineering and electronics from the Institute of Electrical and Electronics Engineers (IEEE) and the Institution of Engineering and Technology.

robustness, we also used DuckDuckGo as another open search engine. When using the general search engines, we restricted the searches to identify PDF filetypes only to extract reports not captured by the other databases while also limiting the number of irrelevant results (recall that ProQuest central includes news and other more general articles).

To search the above databases, the following search query (or a variation of it)⁷ was used:
(tiabsu⁸((‘app’ OR ‘app store’ OR ‘web app’ OR ‘web app store’)

AND

(‘gaming’ OR ‘wearable tech*’ OR ‘smart wearable’ OR ‘voice assistant’ OR ‘smart tv’ OR ‘TV’ OR ‘television’ OR ‘alexa’ OR ‘skills’)

AND

(‘threat’ OR ‘risk’ OR ‘breach’ OR ‘hack’ OR ‘malicious’ OR ‘malware’ OR ‘Trojan’ OR ‘ransomware’ OR ‘spyware’ OR ‘worm’ OR ‘virus’ OR ‘vulnerability’ OR ‘security’ OR ‘protection’ OR ‘privacy’ OR ‘data privacy’ OR ‘verification’)

AND

(‘framework’ OR ‘guideline’ OR ‘recommendation’ OR ‘app development’ OR ‘regulation’ OR ‘review’ OR ‘security check’ OR ‘permission’)

NOT

‘mobile’ NOT ‘IOS’ NOT ‘Android’ NOT ‘google play’ NOT ‘amazon app store’ NOT ‘window mobile app store’ NOT ‘window mobile app’ NOT ‘apple app store’)

The latter (‘NOT’) terms were included to exclude literature concerned with mobile applications as these were already covered by Furnell’s (2021) review.

An academic librarian with expertise in the conduct of systematic reviews was consulted to validate the databases and search terms used. Snowballing (i.e. identifying additional papers cited in articles located through the database searches) was used to identify further relevant papers. However, most of these were excluded due to their publication date (see below) or because they were out of scope (e.g., they were related to a web application).

1.3.2 Study Inclusion Criteria

Studies or reports employing any types of study designs (e.g. qualitative and quantitative including systematic reviews and meta-analyses, Randomised Controlled Trials, cohort studies, case-control studies, cross-sectional surveys, case reports, position papers) were included. All types of information sources were included with the exception of articles that were not available in English or that had to be purchased. Only papers published in or after 2017 were included to ensure their relevancy to current online environments.

For the academic review, we used Endnote to manage the literature database and EPPI Centre

⁷ The search query provided was applied to the ProQuest Central database. Variants of this were applied to the other databases searched.

⁸ The term “tiabsu” indicates the search will be applied to the “title and abstract” of every article searched.

Reviewer (Thomas et al., 2010) to remove duplicates and for the screening and extraction of data. For the purposes of clarity, Figure 1 summarises the decision process used to select articles for the review.




Figure 1. Decision tree used to identify articles for the review.

A summary of the volume of articles identified (and excluded) at each stage of the search process is provided in the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA)⁹ chart shown in Figure 2.

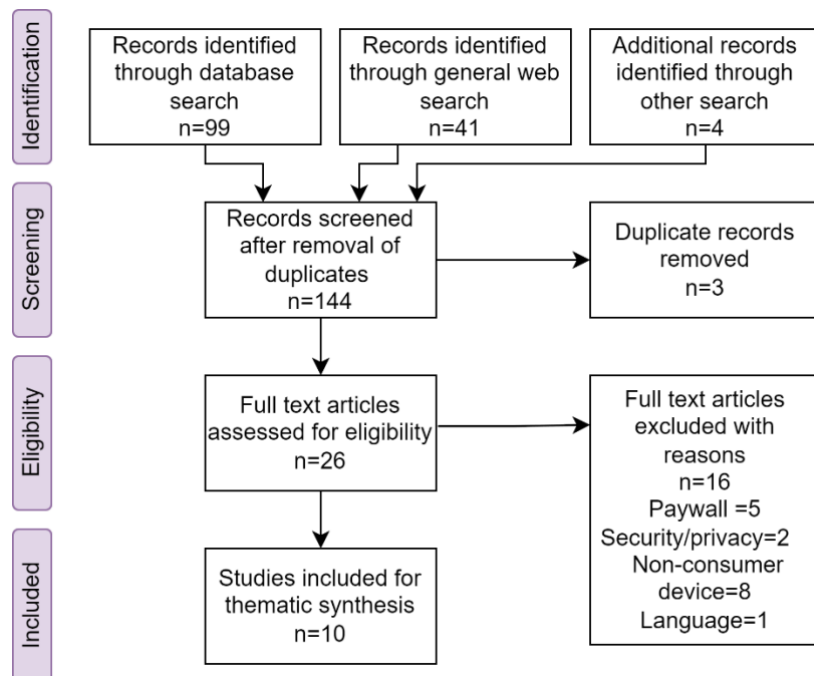


Figure 2. PRISMA chart of the number of articles identified, screened and ultimately reviewed. Additional records were identified through snowballing.

In the first stage of screening, the titles and abstracts of the 144 articles identified were read and assessed against our inclusion criteria. In terms of quality assurance, two researchers independently screened the titles and abstracts of 10% of the (144) identified papers and assessed whether they met the inclusion criteria. Inter-rater reliability¹⁰ was assessed based on the alignment of the two coders in terms of whether they would include or exclude the articles whilst screening the titles and abstracts of the papers identified. There was perfect agreement for this exercise.. Details of the databases

⁹ An evidence-based minimum set of items (27-item checklist and a 4-phase flow diagram) for reporting in systematic reviews and meta-analyses to facilitate and demonstrate preparation and reporting of a robust protocol for a systematic review (Moher et al., 2015).

¹⁰ Inter-rater reliability (IRR) was assessed based on two coding categories (i.e., inclusion versus exclusion) using the prevalence- and bias-adjusted kappa (PABAK) statistic, which controls for chance agreement. For this review, the PABAK score of 1.0 at the Title/Abstract stage indicated perfect agreement between the two reviewers.

searched, and from which the 144 articles were identified are shown in Table 1.

Table 1. Summary of the volume of articles identified (and excluded) at each stage of the search process.

Database/Method	Items	Duplicates	
			Total
ProQuest Central	91	1	90
IEEE Xplore	9	2	7
ACM digital library	2	0	2
Google Search Engine	40	0	40
DuckDuckGo Search Engine	1	0	1
Additional articles (e.g., via snowballing ¹¹)	4	0	4
Total articles	147	3	144

The full texts of the 26 articles that appeared to meet our inclusion criteria were then read and assessed against our inclusion criteria. Ultimately, ten studies were found to meet the inclusion criteria, and for each of these articles, the following data were extracted:

- Year of study
- Author(s) of study
- Country of study
- Publication type (journal, paper in conference proceedings etc.)
- Data analysed
- Study design (e.g., experimental study, focus group, interviews, Delphi method)
- Brief description of study
- Consumer device type
- Threats identified
- Recommendations presented

A thematic analysis (Thomas and Harden, 2008) was then used to synthesise and identify the security/privacy issues and recommendations in the set of studies included in the review.

1.4 Studies identified

More than half (6/10) of the identified articles that met our inclusion criteria were conference proceedings¹² (Tian et al 2017; Celik et al 2018; Acar et al 2020; Escher et al., 2022; Sha et al., 2020; Huang and Hsu, 2017) and the rest were journal articles (Lindqvist, 2017; Johnson et al., 2020; Shi et al., 2021; Sikder, Babun & Uluagac, 2021).

One study was conducted by researchers based in Taiwan (Huang and Hsu, 2017), one by a team in Germany (Escher et al., 2022), one was written by a Finnish researcher (Lindqvist, 2017), one was a collaboration between researchers in the United Kingdom and Australia (Johnson et al 2020), four publications were by researchers in the USA (Sikder, Babun & Uluagac, 2021; Celik et al 2018) or by researchers in the USA who collaborated with researchers in China (Shi et al., 2021), Italy and Germany (Acar et al 2020) or Samsung (Tian et al 2017) and finally, one article was co-authored by researchers in Canada and the UAE (Shah et al., 2020). As such, few articles were identified in the review and only one was published by UK researchers.

¹¹ The process of identifying additional papers cited in articles located through the database searches (Denscombe, 1997).

¹² Conference proceedings in academia refer to a collection of academic papers published in the context of an academic conference.

Most of the studies (8 out of 10) used data collected by the authors, 6 of these applied experimental methods such as threat modelling (Sikder, Babun & Uluagac, 2021; Shi et al., 2021; Acar et al 2020; Celik et al 2018; Tian et al 2017) or a discrete-choice experimental framework (Johnson et al., 2020), one study collected data through interviews (Escher et al., 2022) and another did so using a survey (Sha et al., 2020). Only one study was a literature review (Huang and Hsu, 2017) and only one study was an analysis of legislation (Lindqvist, 2017).

Three of the ten studies related to wearables (Huang and Hsu, 2017; Shah et. Al., 2020; Shi et al 2021). Of these, one focused on the Oculus Virtual Reality (VR) headset and cardboard Augmented Reality/VR headsets (Shi et al., 2021), while the other two focused on smartwatches/bracelets (Huang and Hsu, 2017; Sha et al 2020). Of the seven studies that focused on smart home systems (Johnson et al, 2020; Escher et al., 2022; Sikder, Babun & Uluagac, 2021; Acar et al 2020; Lindqvist, 2017; Celik et al 2018; Tian et al 2017), one specifically focused on the Amazon Echo Dot smart speaker and the webcam Google Dropcam (Escher et al., 2022), while the other six focused on connected platforms including Samsung SmartThings, Philips HUE, LIFX Smart Light, and Amazon Alexa (Johnson et al, 2020; Sikder, Babun & Uluagac, 2021; Acar et al 2020; Lindqvist, 2017; Celik et al 2018; Tian et al 2017).

3. Recommendations Provided

In this section we discuss the recommendations identified. As our search returned so few articles, we include recommendations that are perhaps secondary to the key aims of the review.

For the App store

Devices that are internet connected have access to data that can be intensely private and can reveal information about when the user is asleep, who is in the house, what door is locked with a pin code, or what is being watched on TV or other media. As far as we are aware, tools to assess IoT privacy are not available on the marketplace. To address this issue, in their study, Celik et al. (2018) developed a tool called SAINT that analyses IoT apps to identify sensitive data flows. The tool uses a static analysis, which means that it looks at the programming code of the app (without executing the program) to check that the code structure adheres to industry standards (Chess & McGraw, 2004) and also "models" how information (e.g. that about sensitive device states, such as whether a door is locked/unlocked) flows to and from an app. Analyses were conducted for three major existing IoT platforms (SmartThings, OpenHAB, and Apple's HomeKit). Using this tool, the authors evaluated 168 official and 62 third-party apps and found that more than half (92 and 46, respectively) exposed at least one piece of sensitive data via the Internet or messaging services. The authors consequently recommend that:

Tools such as SAINT should be implemented by IoT app stores so as to determine if the data flow of an IoT app is deemed malicious/dangerous. In such cases, the app should be rejected from the marketplace or modified by the developer.

Security risks extend to devices *within* a smart system, especially those that have full functional access through overprivileged apps. For example, Tian et al. (2017) found that even if an IoT platform is secure, the challenge of malicious apps remains. The authors scanned 180 apps available in the SmartThings marketplace and found that 17 provided access to IoT devices that would allow a malicious actor to send dangerous commands to the app (running on a user's smart devices). Furthermore, they discovered 27 apps that enabled the sending of sensitive data to, and the receipt of commands from, third-party servers. The issue is that a discrepancy can exist between the privileges a user assumes an app has (and what that means it can do or the data that it can access) based on

what an app developer discloses in the app store, compared to the actual privileges the app does have (and what that means it can do or the data that it can access).

In response, Tian et al. (2017) developed an app *SmartAuth* to compare and identify discrepancies between a given app's permissions (established through program analysis¹³) to the developer's representation of them. The analysis of the latter was conducted using Natural Language Processing¹⁴ so that this could be automated (and the burden on the user minimised) and the app used in practice. Tian et al. (2017) used *SmartAuth* to test 180 apps in Samsung's SmartThings store and reported that 17% of the apps analysed were overprivileged. That is, the app requested access permissions that (the analysis conducted using natural language processing suggested) were not described to the user. The authors also evaluated *SmartAuth*'s impact on user decision-making process when installing IoT apps. To do this, they tested 100 study participants (19 - 41 years, average 25.7 years, 59% M, 41% F) with and without *SmartAuth*. They found that when participants were asked if they wanted to install one of two similar apps, one of which was overprivileged, absent the use of *SmartAuth*, 48% of participants installed the overprivileged app through the platform used. For those that used *SmartAuth*, however, only 16% of participants installed the overprivileged app, suggesting the potential value of such an app. Consequently, the authors recommend that:

Tools like *SmartAuth*, which compare the permissions that app developers state an app requires to those that it does, should be utilized by IoT app marketplaces to provide better user protection.

For the App developer

In an experimental study, Sha et al. (2020) examined user's ability to absorb information regarding cybersecurity and privacy issues provided through standard app permission frameworks for third-party apps¹⁵. Their findings suggest that for those tested, the mechanisms used were ineffective with participants failing to demonstrate an informed understanding of the permission requests that were likely to be associated with the apps considered. Sha et al. (2020) also found that 69.37% respondents hold app development companies responsible for application data breaches. Consequently, they recommend that development companies should put proper measures in place to prevent them. They suggest that the use of app permissions alone is not an effective method of protecting users from malicious activity and recommend that:

Additional methods such as the use of popups or smart visualizations that convey the type of data that is being accessed by an app *prior* to its use should be implemented by app developers for a more comprehensive approach.

¹³ Security policy was extracted from the code of the program and was compared to the app's free-text description.

¹⁴ The application of computational techniques to extract semantics from free-text.

¹⁵ Sha et al. (2020) performed two experiments to study the effectiveness of the permissions framework used for third party applications developed for watchOS which runs on Apple Watch hardware. The first study involved 111 participants who were asked to complete a survey with 6 questions about respondents' perception of security and privacy for the smart apple watch permissions framework (e.g., Do you pay attention to application permission requests for Apple Watch? Do application developers use Apple Watch data being gathered for unethical or malicious purposes?). The second study was an experiment that involved 157 participants, split across two randomly selected groups (n = 88; n=69), who were asked to indicate which permissions – from a list they were presented with – they thought a series of apps would require to function. In the first study, 32% of respondents reported that they “always” paid attention to app permission requests and 35% reported that they usually did so. However, in the second experiment, when asked what the permission requests were likely to be, participants performed at the chance level. Moreover, when participants were asked for which Apps there was the greatest risk of misuse of the data collected, they identified the Apple Music application as one of the highest at risk, expressing concern for this app above other applications that may collect more personal/sensitive information, such as the Microphone, Camera, Photos, and Health apps.

Moreover, Sha et al (2020) recommend that:

App developers should refrain from accruing data that is not needed by the application for routine operation and that data retention should be made clear to users (e.g. duration, and how it is processed) in an accessible manner.

Lindqvist (2018) note that data storage and other routine processes are often outsourced to third parties (data processor) by IoT companies (the data controller). Consequently, they analysed the contractual relationship between *controllers* and *processors* dealing with the IoT, including the rights and obligations associated with the General Data Protection Regulation (GDPR)¹⁶. They report that both controllers and processors are obliged directly by law to ensure the security of personal data processing. When these activities involve personal data, data protection legislation applies. However, the IoT (e.g., Smart Things) brings a new definition to the term “product” in that they are a physical object (the hardware) but also exist in the digital space (software services), both existing in online and offline worlds, making it difficult to define contracts among stakeholders. According to GDPR, it is the controller’s responsibility to ensure security compliance. As noted, however, in reality, controllers (e.g., the manufacturer of an IoT device) often outsource to processors (e.g., Google cloud services), who usually have multiple clients and their own sub-processors (i.e., hardware manufacturer->software developer->service provider). This means that although legislation focuses on the controllers, in practice, contractual terms are often imposed by the processors on the controller (who is their client). This also means that the processor holds (personal) data from different controllers, stored in the same location which can result in a significant security risk. In such cases, despite a liability situation arising from a breach of contract by *the sub-processor*, the controller would be held legally accountable. To address this issue, Lindqvist (2018) recommends that:

Serious attention is paid to well thought-through contracts between data processors and controllers that reflect the legislative requirement and obligation for the data controller.

For the End-User

Policy-based Mobile App for Bystanders and Surrounding IoT devices

Smart home systems are becoming increasingly popular and users can easily setup and control their smart devices by downloading apps directly from the vendor’s app market or by developing their own. While this enhances the functionality of the smart home, it can also expose it to security vulnerabilities. Devices collect and process data, and (increasingly) bystanders have no way of knowing if they have been “sensed” by such a device. The types of data collected by IoT devices varies and can include audio-visual if it is collected from surrounding devices such as voice assistants, AR glasses or smart cars. If they are unaware of its collection, bystanders cannot object to the collection of that data, nor will they be aware of its storage, for how long it will be retained or whether the data will be shared and with whom. This kind of scenario can thus lead to a significant intrusion of privacy, with data collected (without permission) including information on people’s movements, personality, health status, sexuality, or political orientation.

Three articles recommended:

A scanning app as a solution to address the security and privacy concerns associated with the growing ubiquity of smart devices and their seamless integration in the physical environment (Escher et al., 2022; Shi et al., 2021; Sikder, Babun & Uluagac, 2021).

For example, Escher et al. (2022) designed a mobile and smartwatch app intended to be downloaded

¹⁶ Regulation in EU law to protect individuals right to data protection and privacy in the European Union and the European Economic Area (2016).

by the end-user to enable the detection of surrounding IoT devices that may be collecting audio-visual data. Although some commercially available IoT devices have visual (e.g., LEDs) indicators to show that they are recording data, they do not provide bystanders with information about data processing and (possible) privacy implications. The (Android) smartphone app designed by Escher et al. (2022) continuously scans for new signals via Bluetooth (which is commonly used by IoT devices). For each IoT device detected by the transparency app, a local database is queried using the IoT device's ID, which returns its privacy policy as well as the terms of its Data Processor (DP). This privacy information includes the type and name of the device, the recording channel (audio or video), the data retention period, involved third parties, use of a trigger/wake word before recording, and whether the connection to the DP is encrypted. Escher et al., (2022) designed the transparency app to be user friendly by displaying information about the privacy policy in digestible colour coded icons along with textual descriptions to illustrate the privacy grade for each detected device.

In user testing of the app, participants were asked to use it and answer questions related to its usability (e.g., if the privacy rating was intuitive), functionality, and to assess whether they would use the app and if they would recommend it to others. Six out of eight participants reported that they would install the app and recommend it. **A missing feature suggested by participants was the addition of an automatic notification from the app when a particularly invasive device was nearby.** As was the option to display the location of the devices on a map, which the authors suggest could be implemented in the future either through data transmitted through Bluetooth Low Energy technology or through crowdsourcing (similar to Google's traffic maps app, Wizz app).

Security-based Mobile App for Smart Home System Owners

One way that data can be maliciously collected from connected homes is through the passive "sniffing" of the network to extract sensitive information about the users and their activities, thereby invading their privacy. "Sniffing" is a form of wiretapping similar to phone tapping but applied to computer networks. To examine this threat, Acar et al. (2020) modelled a local adversary – inside or near a smart home environment – and showed that they could exploit sensor data captured or "sniffed" from network traffic, even if it is encrypted. This type of attack would enable the offender to detect the state of specific devices within the smart environment (which might include residences, hotel rooms, offices of corporations or government agencies) which would, amongst other things, provide them with an understanding of the routine activities of residents or others who use the location, and whether a location is likely to be occupied at a particular point in time. The authors collected user activity and network data from 10 different users in an emulated smart home environment that used a variety of popular smart home devices. With the collected data, Acar et al. (2020) then applied machine-learning approaches and showed that an adversary could achieve a high degree of accuracy (above 90%) when identifying the state and actions of targeted smart home devices and their users. The authors acknowledge that the attack modelled in this paper was highly orchestrated, with the activity of a single user being modelled with a known set of devices. Nevertheless, to reduce the kinds of (side channel) attacks discussed above (see also, Box 1), Acar et al. (2020) recommend:

The use of "spoofed traffic" (false data packets) in installed apps that use open-source environments to obfuscate user activity

To explain, to protect the end-user from privacy leakage, the installed app would generate specific data packets that would mask the activities of the user from a malicious attacker who is trying to sniff data without disturbing the real device traffic. For example, if the end-user is not at home, the spoofed traffic app could generate false activity to mask their absence. The authors showed that this type of traffic "injection" can be used effectively to hide the state of devices from an adversary - reducing the accuracy of sniffing attacks from about 91% accuracy to about 57%. Acar et al. (2020) propose that

future work is needed on hiding home/destination/source location information.

Box. 1 Six cyber-attacks smart home systems are commonly vulnerable to (adapted from Sikder, Babun & Uluagac (2021) and supplemented by the outcomes of this systematic review)

Impersonation attack – a malicious attempt is made to access recorded legitimate smart home voice commands, or by stealing the true user’s credentials.

False data injection – the smart home system is contaminated with a malicious app with forged data (from sensors or voice commands) to perform malicious activities.

Side-channel attack – design imperfections of a legitimate smart home app allow for activities that can be exploited by malicious actors. For example, an offender might be able infer the routine activities of a home-owner by monitoring the activity of devices in their home.

Denial-of-Service – a malicious smart home app disrupts the normal functionality of a smart device by flooding it with requests.

Triggering malicious app – a malicious app exists in the smart home system that is triggered by a specific smart device action (e.g., turn off/on a light in a specific pattern).

Snooping / Sniffing / Eavesdropping attack – a malicious interception, deletion, or modification of data that is transmitted between two devices through an unsecured network. For example, this type of attack can reveal the contents of confidential communication such as authentication credentials to smart home devices.

Sikder, Babun & Uluagac (2021) consider the problem of securing smart homes from malicious apps. They performed threat modelling for five cyber-attacks that smart home systems are commonly vulnerable to (see Box 1) and designed and tested the effectiveness of a security framework and app, Aegis+, that would work across Smart Home Systems. This uses a Markov Chain-based machine learning technique (essentially it learns what types of activities should follow each other) to understand patterns of normal user activity, and consequently understand what anomalous activity might look like for a particular home. The authors collected data from 20 different users performing typical daily activities over a period of 15 days in three different home layouts. Their application could successfully detect malicious activities across four different platforms including Samsung SmartThings, Philips HUE, LIFX Smart Light, and Amazon Alexa. From the perspective of the user, the Aegis+ app contains a web-based emulator that allows the user to click on any device/sensor within their custom smart home layout to understand the nature of the malicious event. Sikder et al. (2021) highlight the benefits of Aegis+ for vendors, end-users, and developers. For vendors, such an app can detect abnormal activity in a customer’s home and verify whether it is coming from malicious code. For example, a fire alarm might be triggered by false data injection from a malicious app downloaded by a customer that the vendor has detected as malicious. The security application would then provide support by advising the customer to delete the malicious app and reinstall the correct app. For end users, it can identify malicious events and notify them in real time (e.g., an end user may purchase a new smart lock but download a malicious app that enables an attacker to unlock the door by impersonating them) giving them an opportunity to act. For developers, the Aegis+ logic extractor can be used to understand whether the app logic in a smart home system is correct or not and to understand the cause of an event so that necessary action can be taken to enhance the security of

multiple devices within a smart home. Sikder, Babun & Uluagac (2021) thus recommend:

The installation of a context-aware security framework by smart home-owners to protect smart home systems that considers multiple users and multiple platforms.

For the Device manufacturer

Security label

In their study, Johnson et al. (2020) considered the fact that it is currently difficult for consumers to obtain information about the security of IoT devices prior to their purchase (see Blythe et al., 2019), which makes it difficult for them to make informed purchasing decisions. To address this market failure, they examined the potential effect that security labels would have on consumer choice in the context of the IoT. While this study was not concerned with apps or app stores, the findings have relevance for the way in which information about apps and app stores might be conveyed to consumers and hence it is briefly reviewed here.

The authors conducted a series of experiments with 3000 adults to estimate how their decision-making and willingness to pay for an IoT product was affected by different types of security label. Participants were asked to make a series of purchasing decisions for either smart security cameras, Smart TVs, smart thermostats, or a wearable. In each case, they were shown a range of alternative products which varied in terms of functionality, price and whether they had a security label. Across the experiments, three different types of label were tested. The first was a binary “seal of approval” label, the second was an informational label (which highlighted positive security features and a security icon), and the third was a graded security label similar to those used to indicate the energy efficiency of an electronic device. At the end of the experiments, participants were also asked for their views on the security labels that they saw. Statistical analyses suggested that – after accounting for other factors (price and functionality) – participants were more likely to say that they would purchase a device if it carried a security label, but the greatest effect was observed for the informational label tested. For this type of label, participants were found to be willing to pay between 29-40% more of the average cost of the devices tested (security cameras, Smart TVs, wearables, and thermostats). The “Seal of approval” label also had a positive effect on consumer choice for all devices but was the least effective when compared to the other types of security labels tested. Importantly, although (simulated) purchasing behaviour indicated that participants preferred a device with higher security (as indicated by the label), they did not think that such devices were immune from hacking. That is, the data suggest that the use of such a label would not have the unintended effect of providing consumers with a false sense of security.

The authors consequently recommended that:

Manufacturers should use a security label to help consumers navigate the market and know which consumer IoT devices have particular levels of security. Further, that care should be taken in designing such a label as different forms of labels may impact consumer behaviour differently.

As noted, while these findings relate to the security of IoT devices, they might also be applied to the ways in which information is provided about apps and app stores to help consumers make informed choices about them.

Product Design Changes

In addition to controlling smart home devices such as smart speakers, voice commands are increasingly used by wearable devices. For example, the AR/VR headset Oculus Quest allows for voice dictation to control the headset, browse the internet or purchase products. This communication (between the individual and the device) can involve the transfer of sensitive information, such as

credit card numbers and private healthcare/bank transaction information, that could result in severe privacy leakage if not secured. Moreover, to work properly, these types of devices currently have to be closely mounted on a user's head and press on distinct parts of the face. Shi et al. (2021) demonstrated that a malicious app installed on such headsets could "eavesdrop" to capture facial muscle movements and bone-borne vibrations during AR/VR voice communication. This could allow an offender to extract sensitive biometric and speech data such as the speaker's gender, and speech content to infer their identity. Unlike their mobile app counterparts, the sensors built into such headsets (accelerometers and gyroscopes) are zero-permission, meaning that they do not require the user to grant access for their use. This is predominantly because motion sensors are necessary for the functionality of the VR headset to simulate (for example) the user's movements in the virtual environment by tracking their head motion. Shi et al. (2021) recommend that manufacturers design-in solutions to prevent these types of eavesdropping attacks on VR/AR headsets and suggest that:

This could be achieved by the addition of sensory "noises" to prevent the reconstruction of facial movements/bone-borne vibrations.

The authors suggest that this could be done by injecting "noises" to the motion sensor readings to perturb the correlation with the speech-associated facial dynamic, without interfering with the AR/VR motion tracking. However, the authors do not specify more precisely how this could be achieved.

They also suggest:

The vendors of AR/VR headsets should add ductile materials between the headset and the user's face to weaken the facial vibrations that would be captured by the built-in sensors (accelerometer/gyroscope), as a thicker headset was found to exhibit a much lower success rate for eavesdropping spoken content (using bone borne vibrations)¹⁷.

3. Conclusion

Across the ten articles synthesized, we identified a number of different threats for which recommendations were presented. These are summarised in Table 2, along with details of the articles in which they were discussed, for those readers who would like to know more about the threats or read more detail about the recommendations identified.

In summary, the recommendations identified included a security framework installed in smart home systems for the detection of malicious apps, and the use of an app that implements "spoofed traffic" to obfuscate the smart home user's network traffic. The use of a mobile phone application that would help users to identify nearby devices that might violate their privacy in some way was proposed, and it was suggested that this mobile application should contain a feature that automatically alerts users when a particularly invasive device was nearby. Although the authors of some of the studies reviewed here demonstrated that users are receptive to downloading mobile apps of this kind as a security solution, it must be noted that care should be taken to ensure that these apps do not introduce a new vector that a malicious actor could exploit.

Product design changes were recommended for AR/VR headsets to prevent users' biometric data from being leaked. It was also suggested that app developers should refrain from accruing data other than that which is required for the routine operation of devices and that data retention should be made clear to users (e.g. duration, and how it is processed) in an accessible manner. It was suggested that

¹⁷ Shi et al. (2021) discovered in their research that the HTC Vive headset had a much lower success rate for eavesdropping (i.e. inferring) words, compared to the Oculus Quest. They attributed this to the fact that a thicker face cover is used with the HTC Vive than the Oculus Quest.

the latter could be achieved through visual aids regarding app permissions as current approaches fail to empower users with an appropriate understanding of the access apps are granted. Research on IoT security labelling suggests that care should be taken when designing such labels/visual aids as different designs can be more or less effective. It was also suggested that app stores could employ a vetting tool that identifies apps with sensitive data flows and checks for variation in the permissions that apps require and those that users are told about. Finally, and from a legal perspective, it was suggested that providers of connected devices need to pay serious attention to the contracts they have with data processors to reflect the rights and obligations brought by the GDPR that regulates 'the rights of the data subjects'.

As a final comment, despite conducting a systematic search of the literature using an extensive array of database search engines, only ten articles were identified that met our inclusion criteria. Moreover, some of these articles were focused on topics that were not the primary focus of this review. This suggests a gap in the academic and related literature regarding recommendations to improve privacy and security for non-mobile apps and app stores that should be addressed.

Table 2. Summary of identified threats (see Box 1) and provided recommendations (Section 2) in this systematic review.

Study	Recommendations	Threats										
		Unauthorised bystander collection of Personal data (Escher et al., 2022)	Hacker identifying lucrative burglary spots (Escher et al., 2022; Johnson et al 2020)	Hackers steal user health data/ cyberattacks e.g. man-in-the-middle attack (Huang and Hsu, 2017; Johnson et al 2020; Sha et al., 2020)	Third party use of personal user data (Huang and Hsu, 2017; Sha et al 2020)	Impersonation attack (Sikder, Babun & Uluagac, 2021)	False data injection (Sikder, Babun & Uluagac, 2021)	Side channel attack (Sikder, Babun & Uluagac, 2021)	Denial-of-Service (Sikder, Babun & Uluagac, 2021)	Triggering malicious app (Sikder, Babun & Uluagac, 2021)	Snooping / eavesdropping attack (Shi et al 2021)	Over-privileged apps with smart device functionality (Tien et al 2017)
Sha et al 2020	Visual aids for app permissions				X							
Sha et al 2020	Limit and make data retention clear				X							
Sikder, Babun & Uluagac 2021	AEGIS+ app					X	X	X	X	X		
Johnson et al 2020	Security label		X	X								
Escher et al., 2022	Bystander app	X	X									
Shi et al 2021	Product Re-design										X	
Lindqvist, 2018	Create clear contracts between data controllers and processors		X									
Acar et al 2020	Spoofed traffic to hide smart device states		X								X	
Tien et al 2017	SmartAuth app											X
Celik et al 2018	SAINT tool for IoT apps			X	X		X					X

References

- Anthi, E., Williams, L., Słowińska, M., Theodorakopoulos, G., & Burnap, P. (2019). A supervised intrusion detection system for smart home IoT devices. *IEEE Internet of Things Journal*, 6(5), 9042-9053.
- Acar, A., Fereidooni, H., Abera, T., Sikder, A. K., Miettinen, M., Aksu, H., ... & Uluagac, S. (2020, July). 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 (pp. 207-218).
- Babun, L., Celik, Z. B., McDaniel, P., & Uluagac, A. S. (2019). Real-time analysis of privacy-(un) aware IoT applications. *arXiv preprint*.
- Blythe, J. M., Sombatruang, N., & Johnson, S. D. (2019). What security features and crime prevention advice is communicated in consumer IoT device manuals and support pages? *Journal of Cybersecurity*, 5(1).
- Bronwyn van der Merwe (2018). Moving From Communication To Conversation: The Next Phase Of Personalisation, B & T Weekly. <https://www.bandt.com.au/moving-communication-conversation-next-phase-personalisation/>
- Byars, L. (1991). Strategic management, formulation and implementation—Concepts and cases. New York: HarperCollins.
- Byrt, T., Bishop, J., and Carlin, J. (1993). Bias, prevalence and kappa. *Journal of Clinical Epidemiology*, 46, 423–429.
- Celik, Z. B., Babun, L., Sikder, A. K., Aksu, H., Tan, G., McDaniel, P., & Uluagac, A. S. (2018). Sensitive information tracking in commodity IoT. In *27th USENIX Security Symposium (USENIX Security 18)* (pp. 1687-1704).
- Chess, B., & McGraw, G. (2004). Static analysis for security. *IEEE security & privacy*, 2(6), 76-79.
- Chi, H., Zeng, Q., Du, X., & Yu, J. (2020, June). Cross-app interference threats in smart homes: Categorization, detection and handling. In *2020 50th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)* (pp. 411-423). IEEE.
- Denscombe M (1997). The Good Research Guide. Buckingham, Open University Press.
- Escher, S., Etzrodt, K., Weller, B., Köpsell, S., & Strufe, T. (2022, March). Transparency for Bystanders in IoT regarding audiovisual Recordings. In 2022 IEEE International Conference on Pervasive Computing and Communications Workshops and other Affiliated Events (PerCom Workshops) (pp. 649-654). IEEE.
- Hermans, L.M., Thissen, W.A.H., (2009). Actor analysis methods and their use for public policy analysts. *European Journal of Operational Research*, 196 (2).
- Huang, K.-C. and Hsu, J.-F. (2017) 'Balance between Privacy Protecting and Selling User Data of Wearable Devices', in: Calgary: International Telecommunications Society (ITS) (14th Asia-Pacific Regional Conference of the International Telecommunications Society (ITS): 'Mapping ICT into Transformation for the Next Information Society', Kyoto, Japan, 24th-27th June, 2017). Available at: <http://hdl.handle.net/10419/168490>.
- Jia, Y. J., Chen, Q. A., Wang, S., Rahmati, A., Fernandes, E., Mao, Z. M., ... & Univiersity, S. J. (2017, February). ContextIoT: Towards providing contextual integrity to appified IoT platforms. In *NDSS Symposium* (Vol. 2, No. 2, pp. 2-2).
- Johnson, S. D., Blythe, J. M., Manning, M., & Wong, G. T. (2020). The impact of IoT security labelling on consumer product choice and willingness to pay. *PloS one*, 15(1).
- Kotler, P. (1998). Marketing management—Analysis, planning, implementation, and control (9th ed.). Englewood Cliffs: Prentice-Hall.
- Lindqvist, J. (2018). New challenges to personal data processing agreements: is the GDPR fit to deal with contract, accountability and liability in a world of the Internet of Things? *International journal of law and information technology*, 26(1), 45-63.
- Moher, D., Shamseer, L., Clarke, M., Ghersi, D., Liberati, A., Petticrew, M., ... & Stewart, L. A. (2015). Preferred reporting items for systematic review and meta-analysis protocols (PRISMA-P) 2015 statement. *Systematic reviews*, 4(1), 1-9.
- Mohsin, M., Anwar, Z., Husari, G., Al-Shaer, E., & Rahman, M. A. (2016, October). IoT SAT: A formal framework for security analysis of the internet of things (IoT). In *2016 IEEE conference on communications and network security (CNS)* (pp. 180-188). IEEE.
- Porter, M.E. (1990) The Competitive Advantage of Nations, Free Press, New York, NY.
- Reed, M.S., Graves, A., Dandy, N., Posthumus, H., Hubacek, K., Morris, J., Prell, C., Quinn, C.H., Stringer, L.C., (2009). Who's in and why? A typology of stakeholder analysis methods for natural resource management. *Journal of Environmental Management*, 90 (5).
- Shah, M. U., Rehman, U., Iqbal, F., Wahid, F., Hussain, M., & Arsalan, A. (2020, November). Access Permissions for Apple Watch Applications: A Study on Users' Perceptions. In *2020 International Conference on Communications, Computing, Cybersecurity, and Informatics (CCCI)* (pp. 1-7). IEEE.
- Shi, C., Xu, X., Zhang, T., Walker, P., Wu, Y., Liu, J., ... & Yu, J. (2021, October). Face-Mic: inferring live speech and

speaker identity via subtle facial dynamics captured by AR/VR motion sensors. In Proceedings of the 27th Annual International Conference on Mobile Computing and Networking (pp. 478-490).

Sikder, A. K., Babun, L., & Uluagac, A. S. (2021). Aegis+ A Context-aware Platform-independent Security Framework for Smart Home Systems. *Digital Threats: Research and Practice*, 2(1), 1-33.

Smith, V., Devane, D., Begley, C. M., and Clarke, M. (2011). Methodology in conducting a systematic review of systematic reviews of healthcare interventions. *BMC Medical Research Methodology*, 11 (15), 1-6.

Thomas, J., and Harden, A. (2008). Methods for the thematic synthesis of qualitative research in systematic reviews. *BMC Medical Research Methodology*, 8(45), 1-10.

Tian, Y., Zhang, N., Lin, Y. H., Wang, X., Ur, B., Guo, X., & Tague, P. (2017). SmartAuth:User-Centered Authorization for the Internet of Things. In 26th USENIX Security Symposium (USENIX Security 17) (pp. 361-378).

Thomas, J., Brunton, J., and Graziosi, S. (2010). EPPI-Reviewer 4.0: Software for Research Synthesis. London: EPPI-Centre Software; Social Science Research Unit; UCL Institute of Education.

Tripp, O., Pistoia, M., Fink, S. J., Sridharan, M., & Weisman, O. (2009). TAJ: effective taint analysis of web applications. *ACM Sigplan Notices*, 44(6), 87-97.

Wang, Q., Datta, P., Yang, W., Liu, S., Bates, A., & Gunter, C. A. (2019, November). Charting the attack surface of trigger-action IoT platforms. In *Proceedings of the 2019 ACM SIGSAC conference on computer and communications security* (pp. 1439-1453).

Wang, Q., Hassan, W. U., Bates, A., & Gunter, C. (2018, February). Fear and logging in the internet of things. In *Network and Distributed Systems Symposium*.

Yamauchi, M., Ohsita, Y., Murata, M., Ueda, K., & Kato, Y. (2019, January). Anomaly detection for smart home based on user behavior. In *2019 IEEE International Conference on Consumer Electronics (ICCE)* (pp. 1-6). IEEE.