PARA: Privacy Management and Control in Emerging IoT Ecosystems using Augmented Reality

Carlos Bermejo Fernandez
Hong Kong University of Science and Technology
Hong Kong SAR
cbf@cse.ust.hk

Lik Hang Lee
KAIST
South Korea
likhang.lee@kaist.ac.kr

Petteri Nurmi
University of Helsinki
Finland
petteri.nurmi@helsinki.fi

Pan Hui
Hong Kong University of Science and Technology
University of Helsinki
Hong Kong SAR, Finland
panhui@cse.ust.hk

ABSTRACT

The ubiquity of smart devices, combined with a lack of information about data garnered by them, make privacy a significant challenge for adopting smart devices. Ensuring users can safeguard their privacy without compromising the devices' functionality requires effective yet intuitive ways to manage personal privacy preferences. Current solutions for privacy management are severely lacking as they are ineffective in making users aware of potential privacy risks or how to mitigate them and as they offer limited support for interaction. As our first contribution, we develop a novel AR privacy management interface (PARA) that uses AR visualization to contextualize data disclosure and improve user's perceptions of privacy threats. Besides offering support for enhancing user's privacy perceptions, our interface supports privacy control on compatible devices through privacy-enhancing technologies. As our second contribution, we systematically study factors affecting privacy perceptions and privacy control for two device classes (smart camera and smart speaker) through a user study with $N = 32$ participants. Our results show that PARA's contextualization and visualization of privacy disclosure strongly affect the participants' privacy perceptions. For privacy control, we demonstrate that our prototype improves the participant's capability to identify risks and provides an effective and easy-to-use mechanism for controlling privacy disclosure, in contrast to existing state-of-the-art privacy management interfaces.

CCS CONCEPTS

• **Security and privacy** → *Usability in security and privacy*; • **Human-centered computing** → **Mixed / augmented reality**; *Graphical user interfaces.*

KEYWORDS

privacy, management, control, graphical user interfaces, smart devices, augmented reality, mixed reality, AR-IoT interaction.

ACM Reference Format:
Carlos Bermejo Fernandez, Lik Hang Lee, Petteri Nurmi, and Pan Hui. 2021. PARA: Privacy Management and Control in Emerging IoT Ecosystems using Augmented Reality. In *Proceedings of the 2021 International Conference on Multimodal Interaction (ICMI '21), October 18–22, 2021, Montréal, QC, Canada.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3462244.3479885

ICMI '21, October 18–22, 2021, Montréal, QC, Canada

ACM ISBN 978-1-4503-8481-0/21/10…$15.00
https://doi.org/10.1145/3462244.3479885

1 INTRODUCTION

Increasing availability of smart devices in our everyday environments is offering opportunities for innovative new services and applications that can facilitate our lives but also make users increasingly vulnerable to intrusions of their privacy [12, 46]. Safeguarding user privacy is essential for smart device ecosystems as otherwise people may not adopt them or only use the devices in a limited capacity without taking full advantage of their capabilities. What makes safeguarding user privacy particularly difficult is the fact that users rarely are *aware* of the sensors that surround them or the privacy risks associated with them [31, 39, 43, 46]. The challenge of this task is further exacerbated by a lack of a proper control interface to manage privacy as current solutions offer limited means to contextualize data gathered by the devices [2, 34]. Contextualization is here understood in broad terms, referring to any factor (e.g., type of collected data, location of the smart device) that is related to the process of private information disclosure via smart devices. Indeed, the few examples of management interfaces available on current devices, e.g., large-size devices: tangible buttons, built-in LCDs, cf. small-size devices: with companion apps on smartphones[1,2], do not contextualize the management of privacy preferences or the information that is presented. The failure to contextualize the data disclosure (e.g., the location-data relevance, the purpose of data collection) can significantly impact users' privacy perceptions and lead to insufficient controls on user privacy [32, 37, 43].

We contribute by proposing Privacy Augmented Reality Assistant (PARA), a privacy-preserving assistant driven by AR for smart devices at home. PARA has been designed to contextualize data

[1]HomeKit: https://www.apple.com/hk/en/ios/home/
[2]SmartThings: https://www.smartthings.com/

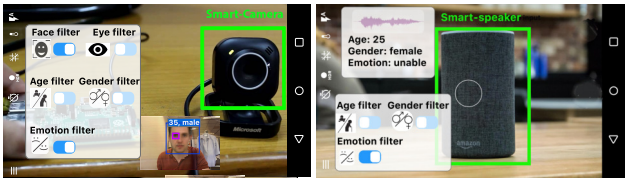

(a) Smart camera. (b) Smart speaker.

Figure 1: PARA for two smart devices. Users point the AR view at a device to see what the collected data and how the privacy filters affect the data.

disclosure and the effects of any actions to control user privacy. PARA thus overcomes the key limitations of existing solutions, and as will be shown, this results in significant improvements in the user's privacy perceptions and control. PARA integrates these functionalities using an AR view on the top of devices. When a user's smartphone points at a smart device, the PARA interface indicates the types of data being collected and enables switching on or off data collection, offering real-time control of privacy management actions. PARA thus allows users to explore in real-time how changes in privacy settings affect data disclosure on devices compatible with our system.

PARA responds to the emerging interaction paradigm of AR with smart devices at home, with the following research questions: *1) How an AR-driven privacy control differentiate from the non-AR counterparts (traditional smartphone UIs and conversational user interfaces)? 2) Does AR-driven user interaction motivate users to manage the privacy settings of smart devices?* PARA features a unified interface [19] for *interacting* with smart devices [3]. We implement an experimental prototype (Figure 1) that supports two highly intrusive yet popular consumer-grade devices: smart cameras (already widely adopted and market is foreseen to expand to 12 billion USD by 2026[3,4] and smart speakers (1 out 10 consumers will own such device [10]). Both devices can expose users to serious privacy leakage via audio and video channels [22, 27, 35].

We conduct a user study with 32 participants to examine the research questions, with emphasis on 1) how different aspects of disclosure context (location, nature of data, purpose) affect user privacy preferences and 2) comparing the AR interface against current privacy management interfaces. We also compare how different configurations of privacy-enhancing technologies affect user's preferences about disclosed data. Our results first highlight that PARA successfully improves the user's perceptions of privacy risks compared to the non-AR counterparts. Our results also show that users with PARA become more aware to the *exact* location of a device and its disclosed data, which result in increased privacy perceptions by 18% from traditional list-based interfaces. In contrast, the non-AR counterparts are not sufficient to motivate users to use or even to be aware of the need for privacy protection.

Taken together, the contributions of this paper are as follows. First, PARA is an AR-based privacy assistant that increases the user's privacy perceptions with a higher intention of applying privacy protection mechanisms. Second, PARA interfaces of privacy filtering demonstrate a higher effect on the willingness of using privacy protection than the existing, non-AR counterparts. Third, our experiments shed light on contextualization of data disclosure and how it can influence the user behaviors to alleviate privacy risks. PARA serves as an intuitive yet informative privacy assistant for smart device ecosystems.

[3]https://www.digitalsignagetoday.com/news/facial-recognition-market-will-reach-12b-by-2026-report-says/

[4]https://www.ifsecglobal.com/video-surveillance/smart-cctv-and-the-internet-of-things-2016-trends-and-predications/

2 RELATED WORK

Privacy management in smart devices. Langheinrich [21] highlights that systems do not seek perfect privacy protection, and illustrates the concept of *Privacy Assistant* to raise the user's privacy awareness. Smart devices do impact on user's perceptions, regardless of owners and non-owners, to the smart devices [28]. Most recent interviews with consumers and experts [12] indicate the urgent need for establishing security and privacy labels for smart home devices. Such labels should inform the users about the sensors, data types, and the granularity of data collection [12]. Colnago et al. [7] used a semi-structured interview to recommend solutions about automation of privacy preferences and notification overloads. Although automation is positively seen by participants as a solution to manage privacy, participants are still concerned about the sources used to generate the automation techniques and request higher flexibility of customizing their privacy preference. The latest work suggests that locator UIs for contextualized images can effectively enhance user's awareness of adjacent smart devices [43]. The authors in [8] propose a smartphone application to manage and control the visual privacy (e.g., facial feature and location) captured by distributed IoT devices. Nevertheless, the privacy management interfaces on smart home devices, e.g., tangible LCD displays and UIs on smartphones, are insufficient for informing users about data collection (UDC), for instance, the location, data type and the targeted use of the shared data [12, 46]. Without appropriate user affordance(s) about user privacy with smart home devices, users could be subject to anxiety from unintended disclosure or even become unwilling to use IoT applications [46].

Privacy threats with smart devices, including wearables, necessitate informing users about their privacy decisions and how device permissions affect the data sensed and shared by these devices. Appropriate visualization is an effective strategy to inform users about the device threats [13, 18]. Additionally, privacy assistant should offer transparent tools to ensure that individual privacy requirements are fulfilled [5]. Some hypothetical scenarios of smart devices [32] trigger the user's privacy reactions to the frequency and types of data collection. PARA is a system solution that leverages interactive AR visualization to inform users about the privacy threats in the user's first person view.

Seizing AR with smart devices. The existing works demonstrate the feasibility of AR for the management of smart devices that primarily own limited form size or (over-)simplified design (i.e. lacking user affordance) [3, 16, 17, 29, 36]. *Smarter Objects* developed by MIT Media Laboratory [16] has been considered the first working prototype to digitally overlay graphical interface on household objects. Various dimensions of AR-IoT have been extended, including

developing scalable spatial registration of IoT devices [17], optimized user interfaces [3], searching and locating the nearby smart devices using the AR Field of View (FOV) [36], visualizing the connections and shared data among smart devices [29], analysis of the effects of AR in privacy perceptions [13]. However, the above works focus exclusively on the technological challenges of AR-IoT interaction but neglect the user perception to smart devices via AR platforms. We develop a novel AR interface for privacy management (motivated by the findings in [13]), and simultaneously evaluate the issues with existing privacy management interfaces.

3 PARA DESIGN AND INTERFACE

In this paper, we primarily focus on the smart home scenario between a device and a user. This section first explains the system design of PARA[5] and the privacy filters offering privacy control. PARA contextualizes the data disclosure to users in real-time and allows users to control their privacy preferences by using privacy filters. Next, we describe the experimental implementation with two smart devices: smart camera and smart speaker (Figure 1).

3.1 System overview

3.1.1 Design principle. Augmented Reality (AR) facilitates natural interaction between users and smart devices [3, 24, 29, 29, 36, 38], and has great potential to improve user privacy [19, 20, 36]. Previous works demonstrate that exposing the physical locations of smart device in the user's vicinity [1, 13, 32, 43] could enhance user privacy. PARA builds on top of previous findings [13] but goes one step further, contextualizing the privacy capability once such threats from nearby smart devices appear. In general, users feel comfortable with smart devices when the shared data is conditioned on informed consent [1]. AR can serve as a privacy-consent interface for user-device interaction [13]. Once the user's device appears in the camera FOV, PARA alleviates the cumbersome and often overwhelming list of data sharing permission [36]. Therefore, we explore the effects on privacy and compare PARA with the widely used approaches on smart devices. In brief, PARA offers an AR interface that enables users to explore and configure the disclosed data in real-time. PARA also contains a software component named *privacy filters* that obfuscates the sensor data before the permission of data disclosure.

3.1.2 System design. PARA consists of three components: **smart devices**, **privacy filters**, and a **user application**. First, the **smart device** transmits the collected 'privacy-protected' data, supported by two components: (i) **privacy manager** that applies the current privacy filter configuration to the *collected data*, and (ii) *databases* storing users' privacy preferences as well as *privacy filter statistics* (e.g., activation and deactivation time with a device). Second, the **privacy filter (PF)** obfuscates the user's personal information such as face, eye, age, gender, and emotion, using data manipulations such as '*deepfake*' approaches [25, 33, 41] (Figure 1). We select these filters following previous works [1, 45] that highlight the risks of video analytics (e.g., emotion recognition [44]) [45] and users concerns in ubiquitous environments regarding the types of collected data [1]. Moreover, these filters are easier for the participants to understand our prototype.

[5]https://solrac1986.github.io/para_smartdevices.github.io/

As shown in Figure 2, when the user enables the *gender and emotion filters* to the smart camera, the system replaces the authentic face by a 'generated' one, which contains a randomly assigned gender, different emotion [33], and similar age to the original [41]. Although there are numerous alternatives to implementing the filters, such as physical filters [4], and middleware [9], we pick embedded filters that reside directly on the devices [42]. Rather than constantly moving a tremendous amount of data to the cloud or edge devices, we perform the inference tasks in the smart device. This approach also reduces privacy threats that might appear during the data transmission or in the cloud [42]. Finally, **user application** displays user data being collected by nearby smart devices and the currently enabled privacy filters. The application detects nearby smart devices using object detection and Bluetooth. Smart devices images are registered in our database. Users via AR can interact with smart devices to alter the privacy filters' status, i.e. enabled/disabled data collection, and inspect how the *privacy filters* affects the collected data in real-time (Figure 2). In scenarios when devices are installed in the same room and the field of view of the camera, the interface will only visualize the location of these devices (motivated by Song *et al.* [43]). The users can select by touching on the screen the device to configure the device and show more related information such as privacy filters and collected data.

3.2 System prototyping and implementation

3.2.1 Prototyping. The prototype of PARA is centered around an Android smartphone (API 14 version) that visualizes the aforementioned system components. Additionally, a Raspberry Pi 3B (Raspberry Pi OS) emulates the connection (Node.js, express 4.1, mongodb 3.2.7, socket 2.0.4) with smart devices (i.e. smart camera and speaker in our experimental setting). The relevant libraries (Tensorflow 2.0, Python 3.6) and machine learning algorithms for **privacy filters** ran on the emulation locally.

User-device interaction. Figure 3 depicts the PARA workflow from the smart device detection to the *privacy filter* updates. We highlight the key steps as follows.

(*1,2*) *Client, smart device's address detection*: the application detects nearby smart devices using object detection. We trained a one-shot learning network [15] in Tensorflow 2.0 to detect the respective smart camera and smart speaker for our study. For occluded objects, the beacon Bluetooth technique can notify the users' device about the proximity of a smart device. We use the similitude technique as various illumination levels and object orientation exist. The client device detects the smart device's IP address using Bluetooth, and it starts a `https` request to access the smart device (*1*). If the connection is successful, the client retrieves the current privacy preference settings from the smart device. The client also receives the current monitored data in real-time (*2*).

(*3, 4*) *Client, privacy filters and data visualization*: PARA transmits the privacy filter status and data visualization through REST architecture[6]. Smart devices are in charge of modifying data visualization and *privacy filters* according to their embedded sensors (*3*). The client displays the device, data, and privacy filters via overlaid Android WebView (*4*). The AR data visualization interfaces employ the web standards (HTML/CSS/JavaScript).

[6]https://socket.io/

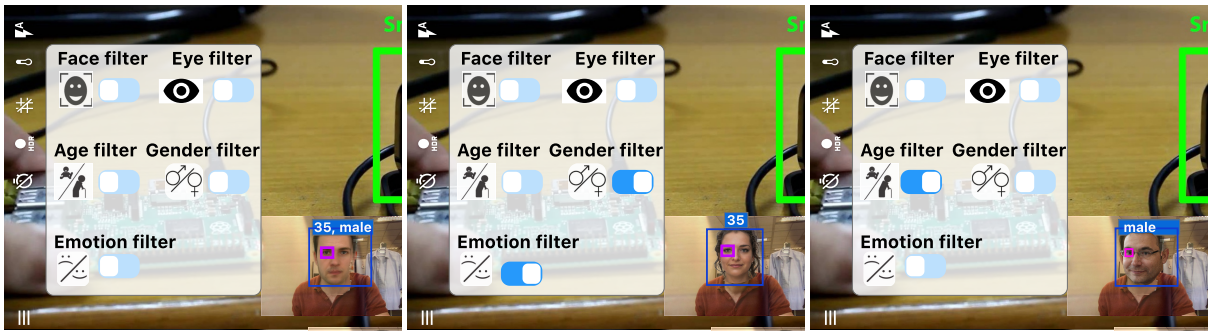

(a) Disabled. (b) Enabled. (c) Enabled.

Figure 2: Privacy filters detailed screenshot when the gender and emotion filter is disabled (2a) and enabled (2b). We can observe the change in the individuals' gender and emotion (smile) while keeping a similar age and displaying the face and eyes. Figure 2c shows the changes when the privacy filter for age is enabled.

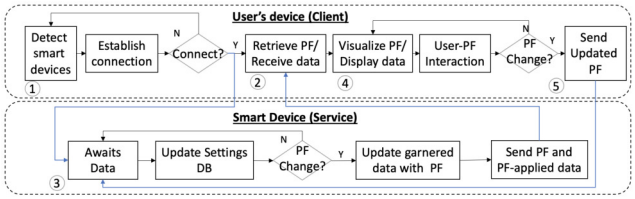

Figure 3: Privacy AR-smart device workflow.

(*5*) *Server, smart device*: the libraries on the smart device apply *privacy filters* on the garnered data (*3*). ML-based **privacy filters** in PARA perform privacy protection of images or audios of the smart camera or speaker, respectively. The smart device processes privacy preferences and user data through a server-side Node.js, where the client application can respond (*5*).

4 DESIGN OF USER STUDY

We implemented two fully working use cases of commercial smart devices: cameras and speakers. News media[7] and studies [27, 32] show that smart cameras and speakers raise most privacy concerns among the device owners [22, 40]. This section first describes the user evaluation design.

4.1 Three experimental configurations

We justify the three configurations and their characteristics, on the basis of contextual information [1, 32] and state-of-the-art privacy managers [8, 14]. The proposed *AR* privacy assistant simultaneously visualizes the contextual information (e.g., the physical location of the smart device) [1, 32] as well as the collected data [8, 32]. In contrast, the non-AR interaction paradigms of *Graphical User Interfaces (GUIs)* and *Conversational User Interfaces (CUIs)* emulates the widely adopted approaches of privacy management with smart devices. For all the configurations, the visual icons that inform users about the collection capabilities (including information inferring) follow [11].

1. Non-AR GUIs. This is analogous to the usual design patterns on mobile and web platforms[8], via the cloud-based smart device access in particular. As shown in Figure 4, the user with GUI privacy assistant receives real-time visualization of the collected data and the possible inferences by third parties [1, 21, 32].

2. Non-AR CUIs. Motivated by privacy bots [14], the user interacts with the smart device using voice commands (Google Assistant). In such an environment, a query-based assistant helps the user to understand and decide their privacy preferences. Figure 4 depicts a conversation between the user and the smart-assistant. The user receives audios indicating each privacy filter and its status (on/off).

3. PARA. As described in the previous section, our proposed AR-mediated privacy assistant enables the user's situational awareness [13, 43] of the smart devices, the real-time visualization of user data (e.g., video, spoken command), and privacy filters (Figure 1).

4.2 Participants and apparatus

4.2.1 Participants. We recruited 32 volunteers (19 male and 13 female) with a mean of 32 years (SD: 8) around the university campus. Their professions are as follow: 10 CS researchers, 7 environmental PhD students, 2 biology researchers, 2 biotechnology researchers, 2 biotechnology PhD students, 3 CS PhD students, and the rest have different backgrounds, such as UX designer or entrepreneur. 68.8% of the participants have some experience with AR applications such as gaming (e.g., Pokemon GO) and 56.2% of the participants have no experience with smart devices.

4.2.2 Apparatus. We use an Android OnePlus 3T device to run the client applications (i.e., through-the-screen AR application). We built our conversational smart assistant using Google 'Dialogflow' and 'Actions on Google'[8]. We use a Raspberry Pi 3B with Wi-Fi capabilities to provide the Node.js server for emulating the smart devices (Raspberry Pi OS, Tensorflow 2.0, express 4.1, mongodb 3.2.7, socket 2.0.4). We conducted our experiment in a quiet office room inside our university campus. We used an iPhone 8 to perform audio recording of participants' answers during the study.

[7] https://www.wired.com/story/the-alexa-amazon-eavesdropping-situation/

[8] https://developers.google.com/actions/

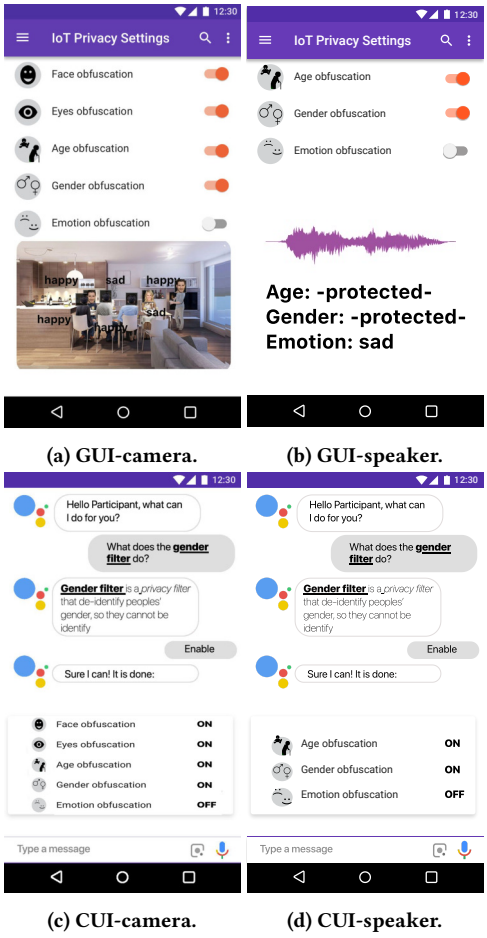

(a) GUI-camera. (b) GUI-speaker.

(c) CUI-camera. (d) CUI-speaker.

Figure 4: Experimental nterfaces for non-AR interaction paradigm of GUI (4a where the users can control the privacy filters (using the corresponding switch button), 4b) and CUI (4c, 4d), where users interact via solely voice-commands with the smart-assistant (no text entry is allowed).

4.2.3 Analysis. For analysing the responses, the data satisfied the assumptions of normality and homogeneity to apply repeated measures ANOVA. We follow Cohen's convention for large effects (range from 0 to 1) [6].

4.3 Task and procedure

4.3.1 Task. The participants use the above interfaces to configure the privacy filters and thus configure the data disclosure. The privacy filters show participants the risks of sharing the collected data with third parties. During our study, the participants are not allowed to turn the devices off or manipulate the sensors. Accordingly, we analyze the participants responses for three experimental configurations and smart devices. The participants are divided into two groups for two experimental use cases: the camera and speaker, in order to alleviate the possible learning effect among smart devices and the corresponding visuals in **privacy filters**. Therefore, we counterbalanced the assignment of participants to each smart device in two equally sized groups (smart camera: $N = 16$; smart speaker: $N = 16$). We first evaluate the *effects of three configurations on users' privacy perceptions*. We compare participants' privacy perceptions (comfort levels), following prior studies [1, 23, 32]. Then, the participants configure their *ideal privacy filters* for each of the counterbalanced assigned configurations.

Table 1: Privacy filter groups.

Group	Description
All	all privacy filter are enabled
None	no privacy filter is enabled
Partial	random selected privacy filters are enabled

4.3.2 Procedure. Participants provided informed consent to participate in this study and be video recorded (by the smart camera). The study was carried out following the General Data Protection Regulation (GDPR) and the IRB regulations of our university. We informed the participants that the experimental data will be de-identified, and all recorded data will be password protected. Afterward, the participants ran through the following procedures.

A. Privacy perceptions. With the counterbalanced order of smart devices and privacy filters, the participants were asked to evaluate their comfort levels according to the smart device, configuration (Non-AR GUI or CUI or PARA), and privacy filter group.

B. Privacy control. The order of the three interaction paradigms and privacy filters are counterbalanced among the participants. We asked the participants to configure their '*ideal*' privacy filter settings, i.e., which filters are enabled/disabled for each configuration to evaluate the *privacy control*.

C. Participants' feedback. Next, the participants were asked a questionnaire regarding the qualitative feedback of each configuration using the technology acceptance model (TAM).

D. Privacy concerns and demographics. We evaluate the participants' general privacy concerns in smart environments using the mentioned Internet users' information privacy concerns (IUIPC) [26]. Finally, we asked the participants five demographic questions: gender, age, profession, whether they have had previous AR experience, and previous experience with smart devices.

The total duration of the procedures is ranged from 30 to 40 minutes for every participants. Participants were rewarded with foods, snacks, sweets, and soft drinks after completing the experiment.

4.4 Ecological validity and limitations

The study setup considers a realistic environment for studying privacy, unlike the online surveys that have been traditionally used to explore privacy [2, 32]. According to the early adoption paradox, consumers with the highest willingness of using IoT technologies have the highest privacy risks in other online activities[9]. Nevertheless, our results must be considered in the context of the limited number of participants recruited with some experience with: smart devices (43.8% of the participants) and smart wearables, e.g., smartwatch (28%), or users/owners of smart devices (28%).

[9] https://blog.f-secure.com/privacy-concerns-cooling-iot-adoption-us-europe/

In terms of limitations, the use of the same physical environment can result in carry-over effects and reduce the accuracy of the privacy perceptions. At the same time, it is essential to offer a realistic environment for interaction to ensure the users can properly contextualize the privacy risks. The study design was designed as a balance of these two factors, and we use counterbalancing to alleviate the carry-over effects. Online surveys used in prior works [1, 32] cannot isolate the effects of contextual factors as they rely on generic scenarios rather than offer a way to contextualize disclosure accurately. To further enhance the level of realism, we limit the study on the analysis of participants' comfort levels according to the presented scenario.

5 EVALUATION RESULTS

According to our user study with 32 participants, we systematically evaluate how the disclosure of contextual information impacts the user perception on both privacy perception and privacy control, and demonstrate how PARA improves these facets compared to existing solutions. We also collected the participants' feedback and their privacy concerns.

5.1 Privacy perceptions

5.1.1 Study design. A factorial study design that assesses users' privacy perceptions according to the configuration used: 2 (between-subjects: smart device types) × 3 (within-subjects: experimental configurations) × 3 (within-subjects: privacy filter group: all, none, and partial enabled filters, see Table 1) factorial design.

5.1.2 Evaluation metrics. We measure participants' privacy perceptions using their comfort levels according to the smart device, experimental configuration, and privacy filter configuration. We follow previous works that study privacy perceptions using comfort levels [1, 32], where low comfort levels (5-point scale) related to participants' low privacy perceptions:

- *Q1.* How comfortable would you feel about using this smart device? (Answer: 1. very uncomfortable - 5. very comfortable)
- *Q2.* Additional comments (Answer: open-ended)

The configurations of privacy filters follow these three groups: *(i) all, (ii) none*, and *(iii) partial*, see Table 1. The participants were instructed to enable/disable the corresponding privacy filters using each configuration of interaction paradigms. Any change in the privacy filter configuration is visualized in real-time. We select five privacy filters motivated by previous works [11, 45]. In [11], the authors describe the capabilities of smart cameras and speakers to infer demographic characteristics of individuals (e.g., gender, age). Moreover, individuals lack awareness regarding the capabilities of such smart devices to infer biological states such as emotions [45]. Therefore, we selected these five privacy filters as representative use cases for our study.

Results. *AR and privacy perceptions:* We first demonstrate that interface type (i.e., configuration) has a significant impact on user's privacy perceptions with PARA having the largest overall impact. Figure 5a depicts the participants' average comfort levels and the distributions of their responses for each experimental configuration. Mixed ANOVA shows that there is a significant effect of the configurations on the participants' comfort levels ($F(2, 60) = 10.84, p < .001, adj. - r^2 = 0.57;\ d = 1.67$), but no the interaction between configurations and smart devices ($F(1, 30) = 6.23, p = 0.32$). The effect size ($d = 1.67$) was found to exceed Cohen's convention for a large effect ($d = .80$). A post-hoc Tukey evaluation shows that the PARA ($p < .05$) influences participants in their perceived value of risks compared to GUI and CUI configurations (see Table 2 for more details), showing that the AR-based interface significantly increases user's perceptions of privacy risks, and consequently reduces their comfort about data disclosure. Our results thus show that contextualization of data disclosure, as supported by PARA, has a significant impact on user's privacy perceptions regardless of device type, whereas existing solutions have much smaller effect.

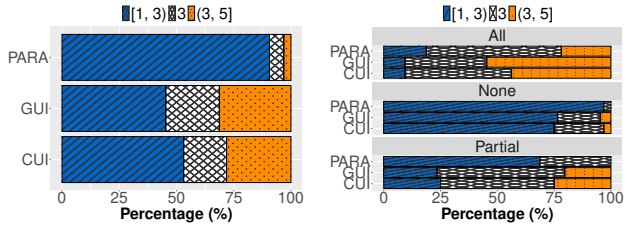

(a) Configurations. **(b) Configurations & filters.**

Figure 5: Privacy perception distributions according to the privacy filter groups: all; none; and partial; and experiment configuration used: GUI, CUI, and PARA. We group the comfort levels according to the following labels: [1,3), levels between one and two; 3, participants responses with comfort level 3 (neutral point); and (3,5], levels between 4 and 5.

Privacy filters and privacy perceptions: In Figure 5, we can observe the valuations of risks according to participants' comfort levels. Mixed ANOVA demonstrates statistical significance in the effect of the privacy filters to the participants' comfort levels ($F(2, 60) = 96.74, p < .001, adj. - r^2 = 0.53$). A post-hoc Tukey evaluation indicates that the enabled privacy filters influence participants in their perceived value of risks ($p < .05$), where we have the lowest perception when no filters (i.e., *none*) are enabled ($M = 1.71,\ 95\% : CI\ [1.56, 1.85]$) in comparison with *all* ($M = 3.30,\ 95\% : CI\ [3.11, 3.49]$) and *partial* ($M = 2.67,\ 95\% : CI\ [2.49, 2.84]$). Our results show a direct relationship between the number of enabled privacy filters and the participants' comfort levels.

AR versus traditional GUI: The main difference between these two interface configurations is the lack of exact location visualization of the smart device in GUI interface. We omit the configuration of CUI as voice-based interaction provides limited user data contextualization (i.e., no visualization of the collected data). The conversational interfaces may mistakenly increase the participant's comfort levels due to its information representation (Figure 4c and 4d). Mixed ANOVA shows significant effect of the contextualization of the location on the participants' comfort levels ($F(1, 30) = 27.64, p < .001, adj. - r^2 = 0.27;\ d = 0.94$), but no the interaction between configurations and smart devices ($F(1, 30) = 4.23, p = 0.4$). The effect size for this analysis ($d = 0.94$) was found to exceed Cohen's (1988) convention for a large effect ($d = .80$). A post-hoc Tukey analysis

Table 2: Summary of the results from privacy perceptions, privacy control, and qualitative feedback of the participants. We mark in bold the statistically significant interaction paradigms.

	Privacy perceptions		Privacy control		Qualitative feedback					
					PU		PEU		IOU	
Interaction	**Mean**	**95% CI**	**Mean**	**95% CI**	**Mean**	**95% CI**	**Mean**	**95% CI**	**Mean**	**95% CI**
GUI	2.71	[2.19, 2.98]	0.7	[0.61, 0.79]	3.34	[3.11, 3.58]	3.79	[3.60, 3.98]	3.22	[2.94, 3.49]
CUI	3.21	[2.34, 2.70]	0.62	[0.52, 0.73]	2.22	[1.80, 2.64]	2.38	[1.99, 2.76]	2.22	[1.88, 2.56]
PARA	**2.29**	[2.03, 2.56]	**0.81**	[0.67, 0.95]	**4.14**	[3.61, 4.67]	**4.16**	[3.38, 4.48]	**3.92**	[3.47, 4.38]

shows that the PARA visualization of the location of the smart device has lower comfort levels ($M = 2.17,\ 95\%\ : CI\ [2.00, 2.33]$) than GUI (data collected visualization) with no location visualization ($M = 2.65,\ 95\%\ : CI\ [2.45, 2.84]$). Participants lowered their comfort levels when the location of the device is displayed in PARA.

Novelty of the use of AR: Despite the use of technologies (i.e., 43.8% of participants see AR as an unfamiliar technology), the privacy perceptions of the participants (i.e., comfort levels) are not influenced by the use of the proposed AR system (PARA).

5.2 Privacy control

5.2.1 Study design. A factorial study design that assesses users' privacy control according to the configuration used: 2 (between-subjects: smart device types) × 3 (within-subjects: experimental configurations) factorial design.

5.2.2 Evaluation metrics. We evaluate the privacy control by recording the number of enabled/disabled privacy filters for each configuration. The available configurable privacy filters in the smart camera and smart speaker are five and three, respectively.

5.2.3 Results of privacy filter preferences. We compare participants answers for each scenario and their ideal privacy filter settings. Mixed ANOVA shows that there is a significant effect of the configuration used on the percentage of privacy filters enabled, regardless of the smart device types ($F(2, 60) = 4.66, p < .001, adj.-r^2 = 0.47$). A post-hoc Tukey evaluation shows ($p < .05$) that the PARA interface increases the probabilities of that a participant will enable more privacy filters in comparison with GUI and CUI (see Table 2). Figure 6a depicts the participants' ideal privacy filter enabled according to the configuration. The PARA interface thus helps users to be involved in the decision process of privacy configuration, especially when a control channel is available. Participants will not keep the default settings when installing a smart device, providing better privacy management than GUIs and CUIs.

5.3 Qualitative feedback

5.3.1 Evaluation Metrics. We measure the qualitative feedback using the TAM, which includes three criteria: perceived usefulness (PU), perceived ease of use (PEOU), and intention to use (IOU). We follow a 5-point scale to measure the three criteria quantitatively, with additional comments (Answer: open-answer). We use the Cronbach alpha to measure the internal consistency of the technology acceptance model scale.

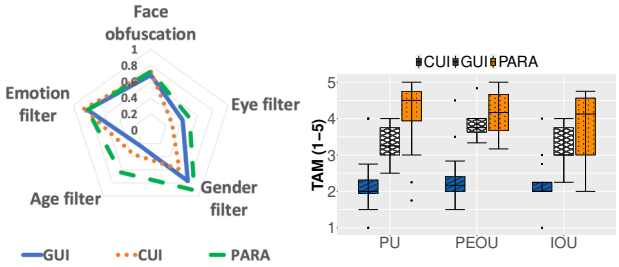

(a) Ideal privacy filter according to the participants and interaction paradigms.

(b) TAM (perceived usefulness, PU; perceived ease of use, PEOU; intention of use, IOU) responses.

Figure 6: Participants' ideal privacy filters and feedback.

5.3.2 Results. We analyze user perceptions of our AR configuration and compare it against the two traditional interaction paradigms (Cronbach's alpha = .73), see Figure 6b and Table 2. Mixed ANOVA shows a significant effect of PARA to the participants' perception of usefulness ($F(2, 30) = 24.78, p < .001, \text{adj.} - r^2 = 0.50$). A post-hoc Tukey test shows that non-AR interfaces (GUI and CUI) and PARA configurations differ significantly ($p < .05$). The participants perceive the usefulness of the proposed PARA configuration higher than GUI and CUI. Mixed ANOVA shows a significant effect of the configurations on the perceived easiness of use ($F(2, 30) = 31.18, p < .001, \text{adj.} - r^2 = 0.62$). A post-hoc Tukey evaluation shows that the PARA is significant higher than CUI configuration, but there is no significant difference with the GUI. Mixed ANOVA ($F(2, 30) = 25.19, p < .001, \text{adj.} - r^2 = 0.51$) shows a significant effect of the configurations on the participants' intention of use. A post-hoc Tukey analysis ($p < .05$)) compares the statistical difference between the PARA configuration which have statistical significance against GUI and CUI. Accordingly, our results highlight how our AR interface is perceived as the most useful and practical. In contrast, the GUI and CUI (primarily text-based metaphor) are perceived as inadequate for managing privacy preferences.

5.3.3 Participants' experiences. The participants mentioned that the CUI configuration is challenging to configure the privacy filters due to the voice-command-based interactions. Some participants were reluctant to use voice interaction, ***P1***: 'voice interaction opens another channel to garner user's information without their awareness.' Participants found the visualization of collected data by GUI and AR are useful during their privacy decision process. Several participants highlighted the difficulties. That is, the current system

might have to visualize the collected data of other types of data (e.g., from light sensors) and how the application of privacy filters might work. Overall, the participants felt that the AR configuration is the 'more fun and easy-to-use', and offers an efficient way to visualize the collected data and device's location.

5.4 Privacy concerns

5.4.1 Evaluation metrics. This metric corresponds with the Internet users' information privacy concerns (IUIPC) [26], which contains three dimensions: collection, control, and perception. We use the questions from [32] as our reference. We use a 5-point scale response for the IUIPC questions.

5.4.2 Analysis and results. The participants are not fully aware of all privacy risks, despite being tech-savvy and having privacy attitudes resembling early adopters. Answers to the IUIPC indicated that the participants shared a common concern that third parties can infer their personal information from their online activities. To analyze the IUIPC, we first performed Principal Component Analysis (PCA) to verify each scale's dimensionality. The PCA showed the original components predicted the total variance: collection (α=0.8), control (α=0.76), and awareness (α=0.68). The participants' concerns are low regarding the collection possibilities of smart devices. Despite the low concerns in data collection environments, our results show the effects of PARA on the participants' perceptions of privacy, where our proposed system lower their comfort levels.

6 DISCUSSION AND LIMITATION

Discussion. The proliferation and heterogeneity of smart devices make efficient user awareness and privacy management challenging. Users with traditional interfaces require a high cognitive workload to get informed about privacy awareness and control, and simultaneously the traditional interfaces do not respond well in the continually shifting context of such environments [30]. Contextualization of disclosure, i.e., collected data and location, becomes a prominent user affordance of evaluating their privacy. Smart devices currently do not provide fine-grained privacy management, no more than the deletion of stored data (e.g., voice logs in Amazon Echo devices) [5]. Our results show that users move beyond the default setting when an appropriate privacy control channel exists, potentially offering privacy filters and thus fine-grained privacy management of smart devices. The individuals using PARA are more sensitive about the *exact* device location and the corresponding data disclosure, and simultaneously being aware of the privacy risks and hence privacy management. Prior works attempted to demonstrate this effect but limited to coarse-grained indicators have failed to show this effect due to being limited to coarse-grained indicators, such as reporting the device's existence in a room with a textual description of the shared data and the management of it [1, 32].

Our evaluation compares the AR-driven privacy assistant with two other interaction paradigms – the graphical user interface (GUI) and the conversational user interface (CUI). AR places intuitive and noticeable privacy visualisations in the blurred boundary between the physical and the digital spatial environment. In other words, AR has been overlaid on the top of the smart devices at home. The visuals directly connect the privacy risk with the devices in the user's private environment, which drives the user toward securing their own privacy proactively. In contrast, GUIs and CUIs offers relatively vague information reflecting the potential privacy leakage [3, 29]. Such visuals cannot sufficiently inform users about the risks and benefits of the collected data and hence are unable to assist the users in their privacy decisions. This makes users more vulnerable to privacy intrusions and reduces the user's trust and comfort to the smart devices especially when intrusions occur.

Limitations. Our evaluation shows that AR results in lower comfort levels than the non-AR counterparts for both the smart camera and speaker. It is worthwhile to mention that the participants show lower comfort levels in general for *all conditions* to smart speakers. Several participants reflected that the news about privacy issues with Amazon Alexa/Echo [22, 27], did impact and bias their privacy awareness to the smart speaker compared to the camera. This also reflects the comfort level is influenced by the user knowledge. Meanwhile, the existing AR metaphor could create excessively visual alerts to the participants. However, the AR design of comforting visuals are out of the study scope in this paper. On the other hand, the current evaluation limits to two noticeable and representative devices. That is, users know their existence. Other studies demonstrated that users have dynamic judgement of data collection across devices and the user's situations [8, 28].

7 CONCLUSION AND FUTURE WORK

This paper presents a fine-grained privacy-enhancing system named PARA, which provides precise contextualization of disclosure and facilitates the user's understanding of their disclosed data. PARA also provides privacy filters to facilitate *privacy control* through AR. Experimental evaluation with 32 participants sheds light on the usage of AR for managing privacy risks associated with consumer-grade smart devices. Information is contextualized and visualized through PARA, providing a significant effect on privacy perceptions. The increased privacy perceptions manifested in a higher desire to control privacy, as evidenced by users enabling a higher number of privacy filters with PARA than traditional interfaces.

For future work, we will extend PARA to other sensing modalities, e.g., magnetometers, light sensors, while the existing scope focuses on visual and audio surveillance through smart cameras and speakers, Also, we will investigate the effect of various AR visualization designs on privacy perceptions. Motivated by the findings in [37], we will display the range of monitoring of different sensors using the AR visualization to study the effects on users' privacy perceptions. Finally, we will consider the usage of PARA in smart public environments, where AR privacy filters could serve multiple users, perhaps with conflicting interests and privacy policy.

8 ACKNOWLEDGEMENTS

We thank our undergraduate researcher assistant for their help during the coding of the participants' comments. This research has been supported in part by project 16214817 from the Research Grants Council of Hong Kong, and the 5GEAR project (Grant No. 318927) and the FIT project (Grant No. 325570) funded by the Academy of Finland.

REFERENCES

[1] Noah Apthorpe, Yan Shvartzshnaider, Arunesh Mathur, Dillon Reisman, and Nick 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 (2018), 1–23.

[2] Paritosh Bahirat, Yangyang He, Abhilash Menon, and Bart Knijnenburg. 2018. A Data-Driven Approach to Developing IoT Privacy-Setting Interfaces. In *23rd International Conference on Intelligent User Interfaces*. ACM, 165–176.

[3] Vincent Becker, Felix Rauchenstein, and Gábor Sörös. 2020. Connecting and Controlling Appliances Through Wearable Augmented Reality. *Augmented Human Research* 5, 1 (2020), 2.

[4] Christopher Champion, Ilesanmi Olade, Constantinos Papangelis, Haining Liang, and Charles Fleming. 2019. The Smart^2 Speaker Blocker: An Open-Source Privacy Filter for Connected Home Speakers. *arXiv preprint arXiv:1901.04879* (2019).

[5] Eugene Cho, S Shyam Sundar, Saeed Abdullah, and Nasim Motalebi. 2020. Will Deleting History Make Alexa More Trustworthy? Effects of Privacy and Content Customization on User Experience of Smart Speakers. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.

[6] J Cohen. 1988. Statistical Power Analysis for the Behavioral Sciences. Hillsdle.

[7] Jessica Colnago, Yuanyuan Feng, Tharangini Palanivel, Sarah Pearman, Megan Ung, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. 2020. Informing the design of a personalized privacy assistant for the internet of things. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.

[8] Anupam Das, Martin Degeling, Daniel Smullen, and Norman Sadeh. 2018. Personalized privacy assistants for the internet of things: providing users with notice and choice. *IEEE Pervasive Computing* 17, 3 (2018), 35–46.

[9] Nigel Davies, Nina Taft, Mahadev Satyanarayanan, Sarah Clinch, and Brandon Amos. 2016. Privacy mediators: Helping IoT cross the chasm. In *Proceedings of the 17th International Workshop on Mobile Computing Systems and Applications*. ACM, 39–44.

[10] Ronan De Renesse. 2017. Virtual digital assistants to overtake world population by 2021. *Ovum, May* 17 (2017).

[11] Serge Egelman, Raghudeep Kannavara, and Richard Chow. 2015. Is this thing on? Crowdsourcing privacy indicators for ubiquitous sensing platforms. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. 1669–1678.

[12] Pardis Emami-Naeini, Yuvraj Agarwal, Lorrie Faith Cranor, and Hanan Hibshi. 2020. Ask the experts: What should be on an IoT privacy and security label?. In *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE, 447–464.

[13] Carlos Bermejo Fernandez, Petteri Nurmi, and Pan Hui. 2021. Seeing is Believing? Effects of Visualization on Smart Device Privacy Perceptions. Proceedings of the 29th ACM International Conference on Multimedia (MM '21).

[14] Hamza Harkous, Kassem Fawaz, Kang G. Shin, and Karl Aberer. 2016. PriBots: Conversational Privacy with Chatbots. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. USENIX Association.

[15] Anfeng He, Chong Luo, Xinmei Tian, and Wenjun Zeng. 2018. A twofold siamese network for real-time object tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 4834–4843.

[16] Valentin Heun, Shunichi Kasahara, and Pattie Maes. 2013. Smarter objects: using AR technology to program physical objects and their interactions. In *CHI'13 Extended Abstracts on Human Factors in Computing Systems*. ACM, 961–966.

[17] Dongsik Jo and Gerard Jounghyun Kim. 2016. ARIoT: scalable augmented reality framework for interacting with Internet of Things appliances everywhere. *IEEE Transactions on Consumer Electronics* 62, 3 (2016), 334–340.

[18] Kosuke Kaneko, Yusuke Tsutsumi, Subodh Sharma, and Yoshihiro Okada. 2020. *PACKUARIUM: Network Packet Visualization Using Mixed Reality for Detecting Bot IoT Device of DDoS Attack*. Springer Science and Business Media Deutschland GmbH, 361–372.

[19] Sarah Krings, Enes Yigitbas, Ivan Jovanovikj, Stefan Sauer, and Gregor Engels. 2020. Development framework for context-aware augmented reality applications. In *Companion Proceedings of the 12th ACM SIGCHI Symposium on Engineering Interactive Computing Systems*. 1–6.

[20] Abhishek Kumar, Tristan Braud, Lik-Hang Lee, and Pan Hui. 2021. Theophany: Multimodal Speech Augmentation in Instantaneous Privacy Channels. Proceedings of the 29th ACM International Conference on Multimedia (MM '21).

[21] Marc Langheinrich. 2002. A privacy awareness system for ubiquitous computing environments. In *international conference on Ubiquitous Computing*. Springer, 237–245.

[22] Josephine Lau, Benjamin Zimmerman, and Florian 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 (2018), 102.

[23] Hosub Lee and Alfred Kobsa. 2016. Understanding user privacy in Internet of Things environments. In *Internet of Things (WF-IoT), 2016 IEEE 3rd World Forum on*. IEEE, 407–412.

[24] Lik-Hang Lee and Pan Hui. 2018. Interaction Methods for Smart Glasses: A Survey. *IEEE Access* 6 (2018), 28712–28732.

[25] Tao Li and Lei Lin. 2019. AnonymousNet: Natural face de-identification with measurable privacy. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*.

[26] Naresh K Malhotra, Sung S Kim, and James Agarwal. 2004. Internet users' information privacy concerns (IUIPC): The construct, the scale, and a causal model. *Information systems research* 15, 4 (2004), 336–355.

[27] Lydia Manikonda, Aditya Deotale, and Subbarao Kambhampati. 2018. What's up with Privacy?: User Preferences and Privacy Concerns in Intelligent Personal Assistants. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*. ACM, 229–235.

[28] Shrirang Mare, Franziska Roesner, and Tadayoshi Kohno. 2020. Smart Devices in Airbnbs: Considering Privacy and Security for both Guests and Hosts. *Proceedings on Privacy Enhancing Technologies* 2020, 2 (2020), 436–458.

[29] Simon Mayer, Yassin N Hassan, and Gábor Sörös. 2014. A magic lens for revealing device interactions in smart environments. In *SIGGRAPH Asia 2014 Mobile Graphics and Interactive Applications*. 1–6.

[30] Vikram Mehta, Daniel Gooch, Arosha Bandara, Blaine Price, and Bashar Nuseibeh. 2021. Privacy Care: A Tangible Interaction Framework for Privacy Management. *ACM Trans. Internet Technol.* 21, 1, Article 25 (Feb. 2021), 32 pages.

[31] M. Mikusz, S. Houben, N. Davies, K. Moessner, and M. Langheinrich. 2018. Raising awareness of IoT sensor deployments. In *Living in the Internet of Things: Cybersecurity of the IoT - 2018*. 1–8.

[32] Pardis Emami Naeini, Sruti Bhagavatula, Hana Habib, Martin Degeling, Lujo Bauer, Lorrie Faith Cranor, and Norman Sadeh. 2017. Privacy Expectations and Preferences in an IoT World. In *Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)*. USENIX Association, 399–412.

[33] Vansh Narula, Kexin Feng, and Theodora Chaspari. 2020. Preserving Privacy in Image-based Emotion Recognition through User Anonymization. In *Proceedings of the 2020 International Conference on Multimodal Interaction*. 452–460.

[34] Katarzyna Olejnik, Italo Dacosta, Joana Soares Machado, Kévin Huguenin, Mohammad Emtiyaz Khan, and Jean-Pierre Hubaux. 2017. Smarper: Context-aware and automatic runtime-permissions for mobile devices. In *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 1058–1076.

[35] Antti Oulasvirta, Aurora Pihlajamaa, Jukka Perkiö, Debarshi Ray, Taneli Vähäkangas, Tero Hasu, Niklas Vainio, and Petri Myllymäki. 2012. Long-term effects of ubiquitous surveillance in the home. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*. ACM, 41–50.

[36] Yongtae Park, Sangki Yun, and Kyu-Han Kim. 2019. When IoT met Augmented Reality: Visualizing the Source of the Wireless Signal in AR View. In *Proceedings of the 17th Annual International Conference on Mobile Systems, Applications, and Services*. ACM, 117–129.

[37] Sarah Prange, Ahmed Shams, Robin Piening, Yomna Abdelrahman, and Florian Alt. 2021. PriView–Exploring Visualisations to Support Users' Privacy Awareness. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–18.

[38] Kirill A. Shatilov, Dimitris Chatzopoulos, Lik-Hang Lee, and Pan Hui. 2021. Emerging ExG-based NUI Inputs in Extended Realities: A Bottom-up Survey. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 11 (2021), 1 – 49.

[39] Peter Shaw, Mateusz Mikusz, Petteri Nurmi, and Nigel Davies. 2019. IoT Maps:Charting the Internet of Things. In *The 20th International Workshop on Mobile Computing Systems and Applications (HotMobile)*. ACM.

[40] Zhiqi Shen, Shaojing Fan, Yongkang Wong, Tian-Tsong Ng, and Mohan Kankanhalli. 2019. Human-imperceptible Privacy Protection Against Machines. In *Proceedings of the 27th ACM International Conference on Multimedia*. 1119–1128.

[41] Sola Shirai and Jacob Whitehill. 2019. Privacy-Preserving Annotation of Face Images through Attc ribute-Preserving Face Synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*.

[42] Mingcong Song, Kan Zhong, Jiaqi Zhang, Yang Hu, Duo Liu, Weigong Zhang, Jing Wang, and Tao Li. 2018. In-situ AI: Towards autonomous and incremental deep learning for IoT systems. In *2018 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. IEEE, 92–103.

[43] Yunpeng Song, Yun Huang, Zhongmin Cai, and Jason I Hong. 2020. I'm All Eyes and Ears: Exploring Effective Locators for Privacy Awareness in IoT Scenarios. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.

[44] Yufeng Yin, Baiyu Huang, Yizhen Wu, and Mohammad Soleymani. 2020. Speaker-invariant adversarial domain adaptation for emotion recognition. In *Proceedings of the 2020 International Conference on Multimodal Interaction*. 481–490.

[45] Shikun Zhang, Yuanyuan Feng, Lujo Bauer, Lorrie Faith Cranor, Anupam Das, and Norman Sadeh. 2021. "Did you know this camera tracks your mood?": Understanding Privacy Expectations and Preferences in the Age of Video Analytics. *Proceedings on Privacy Enhancing Technologies* 2021, 2 (2021), 282–304.

[46] Serena Zheng, Noah Apthorpe, Marshini Chetty, and Nick Feamster. 2018. User Perceptions of Smart Home IoT Privacy. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 200, 200:1–200:20 pages.