<https://helda.helsinki.fi>

Seeing is Believing? Effects of Visualization on Smart Device Privacy Perceptions

Bermejo Fernandez, Carlos

ACM

2021-10

Bermejo Fernandez , C , Nurmi , P & Hui , P 2021 , Seeing is Believing? Effects of Visualization on Smart Device Privacy Perceptions . in MM '21: Proceedings of the 29th
pö ACM International Conference on Multimedia . ACM , pp. 4183 4192 , I
Conference on Distributed Computing Systems , Singapore , 29/11/2020 . <https://doi.org/10.1145/3474085.3475552>

<http://hdl.handle.net/10138/335913>

<https://doi.org/10.1145/3474085.3475552>

acceptedVersion

Downloaded from Helda, University of Helsinki institutional repository.

This is an electronic reprint of the original article.

This reprint may differ from the original in pagination and typographic detail.

Please cite the original version.

Seeing is Believing? Effects of Visualization on Smart Device Privacy Perceptions

Carlos Bermejo Fernandez
HKUST
Hong Kong SAR
cbf@cse.ust.hk

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

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

ABSTRACT

Research on smart device privacy has consistently highlighted how privacy is an important concern for users, but they fail to act on their concerns. While this discrepancy between user perceptions and actions has been consistently reported, currently there is a limited understanding of why this is the case or how the situation can be ameliorated. This paper systematically studies how visualizations in privacy assistants can improve the situation, reporting on two studies that explore the users' privacy perceptions in smart device ecosystems. The first study shows that displaying device location and data type reduces the users' privacy perceptions. Participants also weigh the use of media such as online news as a source to inform users about the possible inferences. The second study analyzes participants' preferences to visualize smart device information and privacy policies using augmented reality. Through these two studies, we derive insights and guidelines on how to design effective privacy assistants and to improve user's knowledge of risks associated with data disclosure in smart home scenarios.

CCS CONCEPTS

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

KEYWORDS

Privacy, Privacy Assistants, Graphical User Interfaces, Augmented Reality, Smart Devices

ACM Reference Format:

Carlos Bermejo Fernandez, Petteri Nurmi, and Pan Hui. 2021. Seeing is Believing? Effects of Visualization on Smart Device Privacy Perceptions. In *Proceedings of the 29th ACM International Conference on Multimedia (MM '21)*, October 20–24, 2021, Virtual Event, China. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3474085.3475552>

1 INTRODUCTION

The ubiquity of smart devices, combined with a lack of information about data collected by them, makes privacy a significant challenge

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MM '21, October 20–24, 2021, Virtual Event, China

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8651-7/21/10...\$15.00

<https://doi.org/10.1145/3474085.3475552>

Figure 1: Example of an interface that visualizes: the collected data (i.e., inferred information and real-time visualization), and exact location of the device.

for emerging smart device ecosystems. Studies on smart device privacy have consistently highlighted how privacy indeed is an important concern for users [33, 43, 56]. At the same time, however, these studies have shown that users fail to act on their concerns and tend to underestimate the risks of being personally exposed to them. While this discrepancy between user perceptions and actions has been consistently highlighted (so-called privacy paradox), currently limited understanding exists of why this is the case or how it can be ameliorated. Indeed, existing research largely has focused on evaluating user's privacy perceptions without systematically analyzing factors that mediate this response [2, 18, 59]. These previous works focus on the privacy expectations that users have regarding smart devices [59], the user's privacy preferences when purchasing one [18], or the lack of full privacy awareness about the capabilities of such devices [2, 59]. Studying privacy perceptions and expectations alone is not sufficient as the capabilities of smart devices are continually evolving, resulting in novel privacy threats that users can struggle to understand [59]. Moreover, how to transmit the information about devices in a manner that users fully understand is a challenging problem due to the complexity of privacy in such ubiquitous environments and the limited capabilities that users can have to fully understand such privacy threats [2]. Indeed, as will be demonstrated in this paper, even technology-savvy users struggle at understanding the sensing capabilities of commercial off-the-shelf smart devices.

Privacy assistants (PA) have been proposed as a potential solution to assist users in managing their privacy in complex and ubiquitous data collection scenarios such as emerging smart device ecosystems [11–13]. PAs can inform users about the data collection practices, e.g., using notifications and allow them to manage privacy via a companion device [13]. Previous work has proposed simple

PAs that are based on traditional interfaces, such as list-based interfaces used by mobile operating systems and shown that they can increase user comfort levels [33, 43]. However, as we will show in this paper, these systems are not effective at contextualizing the disclosure of private information, which results in a situation where the user’s increased comfort is based on false perceptions about the risks associated with the disclosure.

This paper contributes by systematically investigating the broader design space for visualizing contextual information during the management of privacy in smart devices. Specifically, we evaluate how *different visualization interfaces* in PAs influence users’ privacy perceptions and help the users manage the privacy of smart devices. We study visualizations through an augmented reality (AR) interface, shown in Figure 1, that has been motivated by the finding in [55] showing that location visualization increases users’ awareness. Previous works have shown AR to be an intuitive and effective mechanism for supporting interactions with smart devices [8, 23, 25, 42] but prior to our work no effective AR-based privacy assistants for exploring and adjusting privacy preferences have been proposed.

We use the AR interface to analyze users’ privacy perceptions in emerging smart device ecosystems and report on the results from two studies. The first study evaluates the effects of disclosure on ($N = 32$ participants) users’ privacy comfort and investigates how the visualization of different disclosure context factors (device location, device information, and captured data) affects user’s privacy perceptions. To further explore the design space for AR visualization and privacy interfaces, the second study explores design possibilities for an AR-based PA using crowdsourcing. Previous works have shown (remote) crowdsourcing to be an effective and ecologically valid mechanism for eliciting feedback on interface designs [15, 27], offering the possibility to reach larger number of users and being suited to remote interactions. In the study, $N = 100$ participants are shown different AR interface designs that visualize the context factor using simple layouts and color-based schema to describe the privacy settings’ status and we assess the user’s preferences of the different designs.

The results of the studies show that visualizing the *exact* disclosure and location rather is more effective than being informed of them, as is the case with existing PAs. The results also show that participants rarely have detailed knowledge of the true privacy risks associated with smart devices. In fact, the participants often weigh the use of media (e.g., news) to inform them about the possible inferences that third parties can obtain from the collected data. When comparing different interfaces, the results show that participants strongly prefer an AR-based visualization as it provides the most details about disclosure. However, there also is a trade-off between the amount of information that is provided and the time that is needed to make an informed decision. Our paper offers new insights into factors mediating user’s privacy perceptions and provides design guidelines for improving user’s knowledge of risks associated with smart devices using AR-based PAs.

2 RELATED RESEARCH

We organize the related work into three parts: privacy perceptions and individuals’ understandings of privacy and their behavior; privacy assistants for smart home scenarios; and interactions with smart devices using AR.

Privacy perceptions in smart home scenarios. Ubiquitous computing scenarios, such as smart homes, have reached a point where the configuration and interconnections of devices are complex to understand, hampering the user’s mental models of privacy threats [57]. Users generally have minimal understanding of privacy risks [36, 60], due to the limited information about smart devices’ monitoring capabilities [43, 56]. PAs should provide more accurate information to help users with their mental privacy models [2, 36].

Although the vast majority of individuals aim to control the disclosure of their personal information [28], most users do not feel personally targeted from attacks, trusting the companies in charge of their data (e.g., IoT manufacturers) or existing mitigation strategies [57]. Risks are also weighed in comparison with the benefits of using smart devices [30, 61]. Individuals also vary their privacy preferences in data collection environments based on the situation (e.g., who is collecting the data) and on the influence of others, such as friends [32]. Our work is inspired by previous works, especially, the impact of contextual information such as location [55], but with a focus on how these findings can translate into the design of better and more effective privacy interfaces. Indeed, our work examines how different visual interface elements affect and mediate user’s privacy perceptions, improving our collective understanding of how to design effective privacy interfaces.

Privacy assistants and smart devices. Langheinrich [29] highlights that PAs do not seek perfect privacy protection but help to raise privacy awareness. Even when users are privacy-aware and consider the privacy and security of a device before their purchase, there is a lack of information about current privacy and security policies and approaches being used [18]. Users should also be responsible for preventing any privacy violations and understand the business models, but at the same time information about privacy practices should be more accessible to help users [20, 52]. Authors in [14] proposed a PA developed in a distributed privacy infrastructure for visual privacy where users can control what visual information is being collected about them (i.e., facial features, location). Users upload their facial features to a smartphone, and then these devices beacon their privacy preferences in scenarios where there are input video devices. Designing an interface for such PAs can be a challenging task [5]. Colnago et al. [12] used a semi-structured interview to recommend solutions about automation of privacy preferences and notification overloads. The authors found participants to be concerned about the sources for informing and automate the privacy preferences of smart devices. These works have designed prototypes of privacy interfaces without systematically investigating the design space and how different interface elements affect privacy perceptions. Many of these interfaces are also based on traditional designs, such as list-based interfaces used in mobile operating systems, which may increase user comfort but, as we will show, be ineffective at revealing actual threats.

AR interfaces and smart devices. There is limited research providing guidelines to design through-the-screen AR interfaces [15]. There are some prior works on enhancing smart device management using AR that inspire our work. MIT Media Laboratory [23] proposed a system called Smarter Objects to link physical objects with virtual ones. Another example is Reality Editor 2 [41], which allows the users and smart devices to communicate through AR

and Internet. In [25], authors proposed an AR framework to track and interact with smart devices that uses object detection and tracking algorithms to show an improved AR GUI of device features, analogously to [23]. As suggested in [42], AR can provide more natural interactions with smart devices, but it still requires further evaluation to identify the impact on users' interactions with smart devices. Existing interfaces focus solely on interactions with smart objects without examining how to use AR for designing an effective privacy interface. These previous works form a foundation for our visualizations, and our work contributes on top of these by offering insights into how AR can be used for effective PA design.

3 VISUALIZATION AND PRIVACY PERCEPTIONS

We first study how contextualizing the device's information and possible *misuses* of disclosure affect the users' privacy perceptions. Understanding these effects is vital for designing effective privacy assistants that can offer actionable and useful guidance to users. Previous studies have been limited to exploring overall privacy perceptions and comfort levels [33, 43, 57] without providing an understanding of how different privacy related information affects user's privacy perceptions – or their accuracy.

3.1 Study Design

Our study follows a 2 (between subjects: smart device type) $\times$ 3 (within-subjects: interface type) mixed factorial design. We evaluate users' privacy perceptions for two common smart device categories: (a) smart camera and (b) smart speaker and compare three PA interface approaches that differ in the information they offer.

Independent variables. As interfaces we consider three different PA designs: (i) traditional, which is based on list-based menus from smartphones and computer settings; (ii) traditional with additional information, which includes a real-time feed from the camera/speaker and the collected information; and (iii) AR-based interface, which visualizes the location of the smart device and the collected information:

- (1) *List*. In this scenario, we design the list-based interface used in mobile OS and web interfaces [14] (for cloud-based IoT access) to manage applications' privacy (see Figure 2).
- (2) *List+vis*. We enhance the list-based view with real-time visualization of the collected data. We include a real-time visualization of how third parties infer information on the collected data. For example, in the smart camera scenario, users can see how each individual's information is being extracted, such as age, gender, see Figure 2. For example, participants can visualize in real-time what the smart camera is recording (i.e., their faces) and the data inference from the images.
- (3) *List+vis+loc*. In this scenario, the interface shows the collected data and the location of the IoT device. Augmented reality (AR) has been proposed as a strong candidate for offering a unified interface for *interacting* with smart devices [25]. At its best, AR can provide an intuitive way to discover, contextualize, and configure smart devices [7, 24, 42, 48]. The interface provides real-time visualization of data (smart camera: video, smart speaker: spoken command) and the location of the device (see Figure 1). Contrary to other studies that evaluate users'

Figure 2: Traditional interfaces: list and list + visualization.

privacy perceptions in a more general location such as rooms in a home [4, 43], this interface displays the exact location of the device within a room.

We implemented two fully working use cases of commercial smart devices: cameras, and speakers, and let participants interact via the three interface models. Users can detect nearby devices using object detection (one-shot learning) and Bluetooth. We implement ML techniques to extract more granular information (e.g., emotion) from the captured data. This information and the captured data are then sent to the user's device using REST architecture. We focus on these devices as they are representative examples of devices that can expose users to significant risks of unintended data disclosure (i.e., audio and video) and unwarranted surveillance [1, 30, 38, 45, 58, 61]. People generally find microphones and cameras as the two most privacy-invasive devices in smart homes [46]. For these two devices, individuals even sometimes directly control the camera

or microphone privacy by covering or turning them off [30, 46]. Considering these two devices also allows us to compare how user perceptions vary between visual (camera) and non-visual (speaker) modalities.

3.2 Procedure

Effects of visualization. We split the participants into two equal-size groups according to the smart device being tested (smart camera: $N = 16$; smart speaker: $N = 16$). We assigned a counterbalanced interface design to each participant. The participants were asked to evaluate their comfort level according to the smart device and interface used. For each interface, comfort level is evaluated using the following questions:

- Q1. How comfortable would you feel about using this smart device? (Answer: comfort level 5-point Likert scale, 1. very uncomfortable - 5. very comfortable)
- Q2. Additional comments (Answer: open-answer)

The choice of questions was motivated by prior studies on privacy perceptions which have used similar questions [4, 33, 43].

Effects of misuse. We present participants a possible misuse of the collected data using media sources such as articles [10, 40, 49] and research findings [59]. For example: ‘Did you know smart cameras can track your mood?’ based on the findings in [59]. The participants were then asked to re-evaluate their comfort level.

Participant characteristics. We end the study with an evaluation of participants’ privacy concerns and demographic questions. First, participants were asked a short questionnaire about their general privacy concerns in online environments using the internet users’ information privacy concerns (IUIPC) [37], which contains three dimensions (groups): collection, control, and awareness. We use the questions from literature [43] as a reference. Next, we asked participants five demographic questions such as gender, age, profession, whether they have had previous AR experience, and prior experience with smart devices. The background questionnaire was placed at the end of the study to avoid priming the users to be privacy conscious before interacting with the interfaces.

3.3 Apparatus

We use an Android OnePlus 3T device to run the client applications (i.e., through-the-screen AR application). We develop the three applications using the Android SDK 23. For emulating the smart IoT devices (camera and speaker), we use a Raspberry Pi 3B (Raspbian Jessie Linux 4.4.13-v7, ARM Cortex-A53, 1GB RAM) as it has Wi-Fi capabilities to serve the Node.js server. We use the Logitech HD Webcam C270 to emulate a smart camera (plugged in the Raspberry Pi), and a generic speaker to emulate the smart speaker device (we use the microphone of the Logitech webcam as input voice interface). We place the smart devices on a shelf in an office room of our university. We used an iPhone 8 to audio record participants’ answers during the study.

3.4 Participants

We recruited 32 volunteers (19 male and 13 female) using a consecutive sample, ranging from 25 to 44 years (25 to 34: 75%, 35 to 44: 25%) around a university campus. Their professions follow 10 CS

researchers, 7 environmental PhD students, 2 biology researchers, 2 biotechnology researchers, 2 biotechnology PhD students, 3 CS PhD students, 1 business student, 1 UX designer, 1 market and sales, 1 entrepreneur, 1 social worker, and 1 multimedia researcher. 68.8% have some experience with AR scenarios such as gaming (e.g., Pokemon GO, Ingress, or professional related). 56.2% of the participants have no experience with the IoT ecosystem. The other 43.8% have some experience with IoT devices such as work-related (42.9%), home devices (28.5%), or smart-watch (28.6%).

3.5 Ethics

The study was carried out following the General Data Protection Regulation (GDPR) and local IRB regulations. We informed that data would be de-identified, and all recorded data will be password protected and deleted after the study ends. Participants provided informed consent to participate in this study and to be video recorded (by the smart camera).

3.6 Ecological Validity and Limitations

Our study setup considers a realistic environment for studying privacy, unlike online surveys that have been traditionally used to explore privacy [5, 33, 43]. According to the early adoption paradox, consumers most excited about IoT technologies have the highest privacy risks in other online activities [19]. Nevertheless, our results must be considered in the context of the limited number of participants recruited with some experience (43.8% of the participants) with smart devices.

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 control for the carry-over effects. Online surveys used in prior works [4, 33, 43] 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 displayed scenario (e.g., interface, smart device, and location) without further studying what the participants would do with that scenario.

3.7 Analysis Methodology

Statistical analysis. We analyze participants’ comfort levels using mixed and repeated measures ANOVA, as the data satisfy assumptions about homogeneity and normality. We measure the effect using Cohen’s (1988) convention for a large effect ($d = .80$) when the ANOVA test is significant.

Participants’ comments. We follow the methodology in [21] where two independent researchers iteratively coded 30% of the participants’ comments. Then, the researchers met in person to discuss and consolidate the codes from the participants’ comments (inter-annotator agreement 0.78 Cohen’s Kappa). The remaining 70% of the participants’ comments were split equally and coded independently by one of the two researchers.

Figure 3: Summary of the effects of the three interfaces on users' comfort levels (barplots with standard error of the mean).

3.8 Results

Effects of location. Figure 3 illustrates the participants' average comfort levels and the distribution of their answers for the configuration with and without location. One-way repeated measures ANOVA shows a significant effect of the location on participants' comfort levels for the smart camera ($F(1, 15) = 25.70, p < .001, adj. - r^2 = 0.24; d = 1.95$) and smart speaker ($F(1, 15) = 34.28, p < .001, adj. - r^2 = 0.48; d = 0.92$), see Figure 3. Mixed ANOVA shows a significant effect for configuration ($F(1, 30) = 27.64, p < .001, adj. - r^2 = 0.27; d = 0.94$) but no interaction effect between the configuration and the smart devices ($F(1, 30) = 4.23, p = 0.4$). The effect size for this analysis ($d = 0.94$) exceeds Cohen's convention for a large effect ($d = .80$). A post hoc Tukey analysis show that the visualization of the location of the smart device has lower comfort levels ($M = 2.17, 95\% : CI [2.00, 2.33]$) than the interface with no location visualization ($M = 2.65, 95\% : CI [2.45, 2.84]$). The visualization of the additional location information reduces comfort by 18.11% from the interface with collected data visualization only (*List+vis*), while the additional collected data visualization reduces comfort by 14.7% from the reference case, *list*.

Possible misuses. Demonstrating a possible misuse of collected data results in participants raising their privacy perceptions. Statistical analysis indicates that the general comfort levels were significantly lower after we show misuse examples of the garnered data for each IoT device (one-way ANOVA $F(1, 30) = 39.77, p < .001, adj. - r^2 = 0.38; d = 1.02$). The effect size for this analysis, general comfort ($d = 1.02$) was found to exceed Cohen's (1988) convention for a large effect ($d = .80$). According to the information in Table 1, most of the participants were concerned about the location of the device (for good or bad), and only 23% recognize the possible misuse of the garnered data. Our results show that even tech-savvy and privacy-aware are not fully aware of the privacy risks of the collected data even when they know most of an IoT device's monitoring capabilities.

Participants' characteristics. Answers to the background questionnaire indicated that participants shared a common concern about third parties knowing about their personal information inferred from online activities. To analyze the IUIPC, we first performed Principal Component Analysis (PCA) to verify each scale's

Category Tags (usage)		Example
Benefits	location (43%),	P11: "smart speaker: it can help me control my other smart devices"
	health(31%),	P3: "smart camera: it can be useful to see who is at the door from the living room"
Risks	utility(27%)	P6: "smart camera: record me while sleeping"
	location(72%),	P7: "smart speaker: I don't trust the brand due to the reports about collecting data without permission"
	trust(41%),	
	misuse(23%)	

Table 1: Participants' comments during the study.

dimensionality. The IUIPC, PCA showed the three original components predicted the total variance: collection ($\alpha = 0.8$), control ($\alpha = 0.76$), and awareness ($\alpha = 0.68$). Participants' concerns are low when we ask them about the collection possibilities of IoT devices.

Novelty of the use of AR. We compare the interfaces that only vary in one independent variable (visualization of the smart device's location) to evaluate the causation relation between lower comfort levels and AR interface. The use of novel technologies (for some participants (31%) AR is unfamiliar technology) can be also a consequence of such impact on their comfort levels. For such purposes, we further analyzed the responses of participants that are familiar with AR technologies. Mixed ANOVA show that there is a significant effect of the proposed AR interface for participants that are already familiar with AR interfaces ($F(1, 30) = 19.29, p < .001, adj. - r^2 = 0.25; d = 1.53$). Similarly, mixed ANOVA demonstrates a significant effect of the interaction configuration for participants that are familiar with IoT and AR ($F(1, 30) = 14.45, p < .001, adj. - r^2 = 0.28; d = 1.33$), i.e., the changes in privacy perceptions are robust against potential novelty effects.

Summary In line with previous work, the participants had a good, but limited understanding of the data collection capabilities of smart devices. Visualizing the context of data disclosure decreased user's comfort and resulted in more accurate privacy perceptions than the two list-based interfaces. Prior work has shown location and data type to affect privacy perceptions, but ours is the first work to demonstrate that contextualizing these factors through visualization is indeed effective at increasing the accuracy of user's perception. When a misuse of information was presented, all interfaces resulted in a significant drop in comfort, suggesting that initial privacy perceptions can be optimistic and inaccurate. AR generally resulted in lowest comfort across all configurations, suggesting that it can be an effective way to contextualize the disclosure and improve the accuracy of user's privacy perceptions.

4 CROWDSOURCED AR INTERFACE EVALUATION

Our second study further explores the design space for AR visualization of privacy disclosure by evaluating a through-the-screen AR-based PA. Currently there is limited research providing guidelines for through-the-screen AR interfaces [15] and, with this work,

we aim to be the first to give an overview of the users' preferences for using AR-based PA. The results of this study thus supplement our first study which explored the effectiveness of different interface elements. Inspired by prior research on crowdsourcing-based interface evaluation [15, 27, 42], we use Amazon MTurk to quantitatively analyze different interface designs. We develop a simple AR experience survey, where participants are requested to answer two questions about the privacy setting's current status for different interface designs (e.g., color of the icons). On average, the experiment takes 4 minutes (min: 2.2, max: 6, median: 3.4), following similar experiment times as in [15], where the HITs with 10 task takes an average of 9 minutes to complete.

4.1 Study Design

Our study follows a within-subjects design with counterbalanced interface configurations. We vary the independent variables described below for each displayed interface. We focus exclusively on smart cameras to ensure realism. Communicating visualizations of non-visual data, such as audio, can result in misinterpretations and inaccuracies as the situation cannot be accurately contextualized. Our primary goal in this second study is to provide insights on the users' preferences regarding AR interface design rather than explore the data visualizations. We design the interfaces with basic elements to avoid any biases towards the interface elements in the evaluation of preferences (i.e., more rounded buttons, or more subtle overlay menus on the smart device).

Independent variables. We contextualize our experiment using different interface designs used in 2D layouts, following a similar approach as the work in [15]. Motivated by previous studies and features that are familiar with intuitive implications on task performance, the interface was parameterized into design features that correspond to the independent variables of our study:

- *Smart camera video source:* as the contextualization increases users' privacy awareness [4], we display the smart camera feed in the privacy interface. The displayed camera feed image is the same for all participants (same frame and face).
- *Icon-based menu:* motivated by the work [16], we include an icon representation of the collected data type (e.g., eye icon to represent eye detection is enabled).
- *Color-based menu:* we include a two color scheme (red: OFF, green: ON) to highlight each enabled setting's status. Color-based interfaces can effect user's privacy choices [17, 63] as has been shown for users' password creation [17] and privacy awareness in engine search results [63]. We chose colors that were salient and emotionally neutral. Naturally, future versions aimed at larger deployments should improve usability and accessibility.
- *Table layout:* represents each item of the privacy menu (e.g., face, age, gender detectors) in a table layout, following a tile design [50]. Due to the through-the-screen AR interface nature, we include a design with larger icons/text that users can press to enable/disable the data collection.
- *Block sign:* use of the block sign to represent a particular detector is not collecting data. We include a block sign on top of the icons to visualize the status (ON/OFF) of the data collection type.

4.2 Participants

We use the Amazon Mechanical Turk crowdsourcing platform to manage our survey using Human Intelligent Tasks (HITs). We recruit 120 participants using a consecutive sample. The survey was limited to workers that have a 90-100% HIT approval rating, and the number of previous HITs approved greater than 50. We analyze and proofread the answers of the participants before rewarding them. Participants receive a reward of 0.42 USD the day after completing the survey (using the survey code in the MTurk platform). Before posting the survey on MTurk, we ran a pilot study with 20 participants, recruited at our university, to test and improve our survey, such as questions and answer options.

Out of 100 participants, 58 are male, 42 female. The majority of the participants are below 44 (88%). 40% of the participants have a 4-year degree, while 20% have a professional degree. The majority (70%) have good or excellent experience with AR (e.g., AccrossAir, Sun Seeker, AR Doodle). The IUIPC, PCA showed the three original components predicted the total variance: collection ($\alpha = 0.78$), control ($\alpha = 0.63$), and awareness ($\alpha = 0.62$).

4.3 Ethics

We informed that all data would be de-identified, and all recorded data will be password protected and deleted after the study ends. Participants provided informed consent to participate in this study, following the local IRB regulations.

4.4 Procedure

We display to each participant four counterbalanced design interfaces (screenshots). Participants that first electronically agree to the consent are presented with a brief description of the task. Participants are then presented with four different design interfaces (e.g., color-based with icons and table layout menu) and are asked two questions per condition that ask them to enumerate what data is being collected and what not in the least amount of time:

- What data is being collected with the current settings (e.g., face, eye)? (Answer: multiple-option, with the following choices, face, gender, emotion, age, eyes)
- What data is not being collected with the current settings (e.g., face, eye)? (Answer: multiple-option, with the following choices, face, gender, emotion, age, eyes)

These two questions allow us to ensure high-quality participants' responses. At the end of these four conditions, we asked participants to rate the best design (of the four previously presented) on a 5-point Likert scale. Finally, we ask participants some demographic questions, privacy concerns using the IUIPC using a 5-point scale, and experience with AR.

4.5 Ecological Validity and Limitations

Crowdsourcing benefits from reaching a larger pool of participants, while it suffers from a potentially reduced quality of responses. For visualization tasks the main issue is that the visualization may struggle conveying a realistic context, which can result in inaccurate user perceptions [22]. We overcome this issue through the overall design of our two studies. The first study focused on a realistic environment whereas the second study focuses on capturing a

Figure 4: Examples of different interface designs according to the design features.

larger sample of respondents. The elements that are included in the two studies have been tailored to account for potential limitations, and our second study focuses on evaluating specific aspects of AR interface designs rather than evaluating their effectiveness in realistic contexts. Concerns in terms of data quality and reliability are mitigated by (i) the larger sample we have in the second study; (ii) the focus on workers with high score; and (iii) the validation process that was used to verify the validity of the responses. Our first study also serves as reference for validating the responses in the second study. Prior works on crowdsourcing based interface evaluation [15, 22, 27] have shown the relative differences between laboratory studies and crowdsourced studies to be minimal, particularly when the focus is on specific user interface elements.

One limitation of the study is that the participants are unlikely to be representative of the overall population. Prior works have shown privacy perceptions and decisions to be influenced by cultural differences, technical skills and background [2, 9, 35, 47]. In our sample, the participants have higher technical skills and privacy concerns than ordinary citizen [26], and suffer from limited diversity [6] – even if the sample captures respondents from different countries and cultures. Comparing the results between our two studies shows the results to be stable, despite relying on different sampling mechanisms and coming from different populations.

Another limitation relates to the design space that is being evaluated. We considered five independent variables designed to isolate the factors under study and to avoid any possible novelty and side-effects that sophisticated interface layouts and elements could induce. This is also in line with the guidelines proposed by Poupyrev et al. [50] for evaluating AR interface designs. Our results serve to provide insights into the design space of AR interfaces and offer guidelines on how to design AR-based privacy assistants.

4.6 Results

Task completion. The median task completion is 21.87 seconds for all interface designs. We can observe that the design with a color-based menu, no video source, no table layout, and no blocking sign results in the fastest completion time, 12.1 seconds (95% : CI [9.8, 17.57]), see Figure 5a). On average, the design solutions that do not include the video source have a lower completion time of 16.1 seconds (95% : CI [7.69, 24.57]).

Highest rated interface. We asked participants for the best interface design according to their preferences (e.g., P12: “easier to read information”, P53: “likeness”). Figure 5b depicts the highest rated

Figure 5: Resulting interfaces from the crowdsourcing study.

user design according to participants’ preferences, where we can see that despite being slower than the version without the video source from the camera, it is preferred due to the simplistic visualization of the privacy settings (i.e., color-based, icons, table layout). As mentioned by several participants (e.g., P12, P67, P78): “easier to read information” displayed in the interface.

Completion task and preferences. Our results show that the fastest interface to complete the task is not the most preferred one. Although, the interface design that includes the video source has a median completion task time of 23.3 seconds (95% : CI [13.1, 28.76]), where the participants have the highest completion task times with the video source only interface (Figure 4b).

5 DISCUSSION

Design implications. The paper serves as groundwork for the design space of PA interfaces to interact with smart devices’ privacy, where users are not fully aware of the privacy risks. The wearable ecosystem is a good example as it usually involves minimal in-/output interfaces (e.g., small screens, or only buttons and LEDs) and includes a myriad of sensors (e.g., heart-rate monitor) that can infer users’ behavior without their awareness.

- (1) *Visualization of contextual information.* The analysis of the different visualization interfaces shows that AR interfaces overcome many of the limitations that more traditional interfaces have when displaying the contextualized information about smart devices. Interestingly, the effects of location during this process have more impact (18.11% reduction on the comfort level) than only the visualization of the collected data. The use of media as a source to display possible misuse of collected information can help users further understand the privacy risks of such practices. Designers should display the location of the device to fully inform the users when managing their privacy, and in

cases that the inferred data can be invasive, they can opt to use media sources to exemplify the possible threats/risks of data sharing with data collectors.

- (2) *AR interfaces for PAs.* The crowdsourcing evaluation shows how users highlight the importance of a simple design to transmit privacy information as efficiently as possible. Color-based interfaces are preferred by the participants and are the interface with the fastest completion task time. The addition of visual information about data disclosure, such as the camera feed and inferred data from the image reduces task completion time but does so at a cost of interface preference. Designers should evaluate the requirements of their implementation and those of users to implement the appropriate AR-based interface for PAs.

User groups and use cases. Our work focuses on a scenario where an ordinary user is configuring her smart device. For example, when purchasing the device and setting the privacy preferences for the first time or updating the preferences. Another example is managing preferences in environments where someone else has installed the devices, e.g., smart offices or rental homes. The primary focus of our study is on general consumers that have no prior experience with smart devices so the visualization effects are effective to broader demographics. However, our results and tools are equally useful for other groups, such as device administrators, by offering support in finding and managing the data flows of the devices.

Comparison between IoT devices. Participants' answer patterns do not differ according to the device (speaker or camera) showed in the first user study. Comfort levels for both devices are lower when participants use the *List+vis+loc* interface compared with the other interfaces. For the smart speaker case, participants felt lower comfort-levels, due to the current news of privacy issues with Amazon devices [44]; for example, **P23, P7**: "I have read news about the current privacy problems of Amazon speakers." The visualization of the location might remind participants about the privacy leakage news in the media. The results build on top of previous works [43, 55], showing the effects of the smart device's location on users' privacy perceptions.

Privacy concerns and misuse. Our results highlight how awareness of device capabilities is not sufficient for informing users of the associated privacy risks. Previous work has shown that users are generally highly concerned about IoT privacy risks [5, 33, 43], but our results suggest that these concerns do not translate to individual devices that users choose to deploy in the environment. Only once we presented a misuse of information, the user's attitudes and comfort levels changed significantly. Following the guidelines in [12] the sources for information regarding possible misuse should be selected by users using crowdsourced recommendations and authoritative sources.

Design challenges with AR-based interfaces Our study was carried out in a controlled environment to avoid issues related to tracking objects or similar that could hinder the evaluation [62]. In our study, we do not mention the possible privacy threats that can come with the use of AR approaches [3, 31, 51]. Related works used centralized approaches such as edge devices to transmit information between PAs and smart devices [13]. Future AR implementations could take advantage of similar designs to transmit information

between smart devices and PAs, and use location or Bluetooth to discover devices in the vicinity of the user. Moreover, future implementations of AR applications will be integrated into head-mounted displays or smartglasses with the additional challenges of input interactions such as gestures [34, 54].

Privacy in semi-public and public locations Our user study considered representative examples of sensors in a private environment. Other scenarios, such as public environments or organizational settings where users are bystanders and where the locations of sensors are apriori unknown, require further investigation [53].

Future work. Future directions of our work will include a broader design space of current AR-based interfaces and interactions. The current study is the first systematic approach to evaluate the feasibility and effects of AR-based PAs to manage the privacy of smart devices. We aim to study more sophisticated AR-based interfaces such as the integration of the visualization of the coverage radius of microphone and camera for the smart speakers and cameras. For example, PA implementations using smart glasses that allow users to interact with the physical world and provide natural interaction between users and smart devices (e.g., hand gesture) [8, 39].

6 SUMMARY AND CONCLUSION

This paper systematically studied factors affecting users' privacy comfort, and perceptions in smart device ecosystems, and the role different visualizations can play in improving user's comfort and perceptions. The first study considered how prior knowledge and different information (device location, device information, data type) and prior knowledge affect privacy perceptions and how visualization can influence these perceptions. Our results show that knowing the *exact* device location and data type reduces the user's comfort and results in privacy perceptions that are more accurate relative to the actual privacy threat. The effect of visualizations on privacy perceptions is dependent on the user's prior knowledge of devices and the data they collect. Exposing users to a (benign) misuse of information has a significant impact on comfort. Visualizing disclosure and location is more efficient at improving awareness than merely informing the user and contextualizing data disclosure through visualization together with misuse has the highest effect on user perceptions. The results of our second study on crowdsourced interface evaluation show that a combination of color-based schemas, a simple interface, and the video source of the camera is the preferred design for participants, but that there are trade-offs between the needed interaction time and level of visualization. Our study offers new insights into factors mediating users' privacy perceptions, provides design guidelines for improving user's knowledge of risks associated with emerging smart device ecosystems, and offers guidelines on the design of effective AR-based privacy assistants.

7 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] Noura Abdi, Kopo M Ramakapane, and Jose M Such. 2019. More than smart speakers: security and privacy perceptions of smart home personal assistants. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS)*. 2019.
- [2] Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. 2015. Privacy and human behavior in the age of information. *Science* 347, 6221 (2015), 509–514.
- [3] Alessandro Acquisti, R Gross, and F Stutzman. 2011. Privacy in the age of augmented reality. *Proc. National Academy of Sciences* (2011).
- [4] 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.
- [5] 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.
- [6] Christoph Bartneck, Andreas Duenser, Elena Moltchanova, and Karolina Zawieska. 2015. Comparing the similarity of responses received from studies in Amazon's Mechanical Turk to studies conducted online and with direct recruitment. *PLoS one* 10, 4 (2015).
- [7] Sandhya Baskaran and Hari Kumar Nagabushanam. 2018. Relational localization based Augmented reality Interface for IOT applications. In *2018 International Conference on Information and Communication Technology Convergence (ICTC)*. IEEE, 103–106.
- [8] Vincent Becker, Felix Rauchenstein, and Gábor Sörös. 2020. Connecting and Controlling Appliances Through Wearable Augmented Reality. *Augmented Human Research* 5, 1 (2020), 2.
- [9] Steven Bellman, Eric J Johnson, Stephen J Kobrin, and Gerald L Lohse. 2004. International differences in information privacy concerns: A global survey of consumers. *The Information Society* 20, 5 (2004), 313–324.
- [10] Niraj Chokshi. 2019. Is Alexa Listening? Amazon Echo Sent Out Recording of Couple's Conversation. <https://www.nytimes.com/2018/05/25/business/amazon-alexa-conversation-shared-echo.html>. [Online; accessed 16-September-2020].
- [11] Richard Chow. 2017. The Last Mile for IoT Privacy. *IEEE Security & Privacy* 15, 6 (2017), 73–76.
- [12] 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.
- [13] 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.
- [14] Anupam Das, Martin Degeling, Xiaoyou Wang, Junjie Wang, Norman Sadeh, and Mahadev Satyanarayanan. 2017. Assisting Users in a World Full of Cameras: A Privacy-Aware Infrastructure for Computer Vision Applications. In *Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2017 IEEE Conference on. IEEE, 1387–1396.
- [15] John J Dudley, Jason T Jacques, and Per Ola Kristensson. 2019. Crowdsourcing interface feature design with Bayesian optimization. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [16] 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.
- [17] Serge Egelman, Andreas Sotirakopoulos, Ildar Muslukhov, Konstantin Beznosov, and Cormac Herley. 2013. Does my password go up to eleven? The impact of password meters on password selection. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2379–2388.
- [18] Pardis Emami-Naeini, Henry Dixon, Yuvraj Agarwal, and Lorrie Faith Cranor. 2019. Exploring how privacy and security factor into IoT device purchase behavior. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [19] F-Secure. 2018. Privacy Concerns Cooling IoT Adoption in the US and Europe. <https://blog.f-secure.com/privacy-concerns-cooling-iot-adoption-us-europe/>. [Online; accessed 16-September-2020].
- [20] Casey Fiesler and Blake Hallinan. 2018. "We Are the Product": Public Reactions to Online Data Sharing and Privacy Controversies in the Media. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [21] Elmar Hashimov. 2015. Qualitative Data Analysis: A Methods Sourcebook and The Coding Manual for Qualitative Researchers: Matthew B. Miles, A. Michael Huberman, and Johnny Saldaña. Thousand Oaks, CA: SAGE, 2014. 381 pp. Johnny Saldaña. Thousand Oaks, CA: SAGE, 2013. 303 pp.
- [22] Jeffrey Heer and Michael Bostock. 2010. Crowdsourcing graphical perception: using mechanical turk to assess visualization design. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 203–212.
- [23] 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.
- [24] Zhanpeng Huang, Weikai Li, and Pan Hui. 2015. Ubii: Towards seamless interaction between digital and physical worlds. In *Proceedings of the 23rd ACM international conference on Multimedia*. ACM, 341–350.
- [25] 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.
- [26] Ruogu Kang, Stephanie Brown, Laura Dabbish, and Sara Kiesler. 2014. Privacy attitudes of mechanical turk workers and the us public. In *10th Symposium On Usable Privacy and Security (SOUPS)* 2014. 37–49.
- [27] Steven Komarov, Katharina Reinecke, and Krzysztof Z Gajos. 2013. Crowdsourcing performance evaluations of user interfaces. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 207–216.
- [28] Michelle Kwasny, Kelly Caine, Wendy A Rogers, and Arthur D Fisk. 2008. Privacy and technology: folk definitions and perspectives. In *CHI'08 Extended Abstracts on Human Factors in Computing Systems*. 3291–3296.
- [29] Marc Langheinrich. 2002. A privacy awareness system for ubiquitous computing environments. In *international conference on Ubiquitous Computing*. Springer, 237–245.
- [30] 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.
- [31] K. Lebeck, K. Ruth, T. Kohno, and F. Roesner. [n.d.]. Towards Security and Privacy for Multi-User Augmented Reality: Foundations with End Users. In *2018 IEEE Symposium on Security and Privacy (SP)*, Vol. 00. 807–823. <https://doi.org/10.1109/SP.2018.00051>
- [32] Scott Lederer, Jennifer Mankoff, and Anind K Dey. 2003. Who wants to know what when? privacy preference determinants in ubiquitous computing. In *CHI'03 extended abstracts on Human factors in computing systems*. 724–725.
- [33] 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.
- [34] Lik Hang Lee, Tristan Braud, Farshid Hassani Bijarbooneh, and Pan Hui. 2020. UbiPoint: towards non-intrusive mid-air interaction for hardware constrained smart glasses. In *Proceedings of the 11th ACM Multimedia Systems Conference*. 190–201.
- [35] Yao Li, Alfred Kobsa, Bart P Knijnenburg, and MH Carolyn Nguyen. 2017. Cross-cultural privacy prediction. *Proceedings on Privacy Enhancing Technologies* 2017, 2 (2017), 113–132.
- [36] Heather Richter Lipford, Andrew Besmer, and Jason Watson. 2008. Understanding Privacy Settings in Facebook with an Audience View. *UPSEC* 8 (2008), 1–8.
- [37] 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.
- [38] 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.
- [39] 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.
- [40] Cara McGoogan. 2019. Why your smart TV is the perfect way to spy on you. <https://www.telegraph.co.uk/technology/2017/03/08/smart-tv-perfect-way-spy/>. [Online; accessed 16-September-2020].
- [41] MIT media lab. 2019. Reality Editor. <http://realityeditor.org/>. [Online; accessed 16-September-2020].
- [42] Konstantinos Michalakakis, John Aliprantis, and George Caridakis. 2018. Visualizing the internet of things: Naturalizing human-computer interaction by incorporating ar features. *IEEE Consumer Electronics Magazine* 7, 3 (2018), 64–72.
- [43] 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, Santa Clara, CA, 399–412. <https://www.usenix.org/conference/soups2017/technical-sessions/presentation/naeini>
- [44] Lily Hay Newman. 2018. Don't freak out about that Amazon Alexa eavesdropping situation. <https://www.wired.com/story/the-alexa-amazon-eavesdropping-situation/>. [Online; accessed 16-September-2020].
- [45] 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.
- [46] 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 (Pittsburgh, Pennsylvania) (UbiComp '12)*. ACM, New York, NY, USA, 41–50. <https://doi.org/10.1145/2370216.2370224>

- [47] Leysia Palen and Paul Dourish. 2003. Unpacking "privacy" for a networked world. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 129–136.
- [48] 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.
- [49] Ty Pendlebur. 2019. Here's why your next wireless speaker will listen to your every word. <https://www.cnet.com/news/your-next-wireless-speaker-will-listen-to-your-every-word/>. [Online; accessed 16-September-2020].
- [50] Ivan Poupyrev, Desney S Tan, Mark Billinghurst, Hirokazu Kato, Holger Regenbrecht, and Nobuji Tetsutani. 2002. Developing a generic augmented-reality interface. *Computer* 35, 3 (2002), 44–50.
- [51] Franziska Roesner, Tadayoshi Kohno, and David Molnar. 2014. Security and privacy for augmented reality systems. *Commun. ACM* 57, 4 (2014), 88–96.
- [52] Eva-Maria Schomakers, Chantal Lidynia, and Martina Zieffle. 2018. Hidden within a Group of People-Mental Models of Privacy Protection.. In *IoTBDs*. 85–94.
- [53] 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.
- [54] Jiayu Shu, Rui Zheng, and Pan Hui. 2016. Cardea: Context-aware visual privacy protection from pervasive cameras. *arXiv preprint arXiv:1610.00889* (2016).
- [55] 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.
- [56] Eric Zeng, Shrirang Mare, and Franziska Roesner. 2017. End User Security and Privacy Concerns with Smart Homes. In *Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)*. USENIX Association, Santa Clara, CA, 65–80. <https://www.usenix.org/conference/soups2017/technical-sessions/presentation/zeng>
- [57] Eric Zeng, Shrirang Mare, and Franziska Roesner. 2017. End user security and privacy concerns with smart homes. In *Thirteenth Symposium on Usable Privacy and Security ({SOUPS} 2017)*. 65–80.
- [58] Nan Zhang, Xianghang Mi, Xuan Feng, XiaoFeng Wang, Yuan Tian, and Feng Qian. 2019. Dangerous skills: Understanding and mitigating security risks of voice-controlled third-party functions on virtual personal assistant systems. In *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 1381–1396.
- [59] 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.
- [60] Elena Zheleva and Lise Getoor. 2009. To join or not to join: the illusion of privacy in social networks with mixed public and private user profiles. In *Proceedings of the 18th international conference on World wide web*. 531–540.
- [61] 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 (Nov. 2018), 20 pages. <https://doi.org/10.1145/3274469>
- [62] Feng Zhou, Henry Been-Lirn Duh, and Mark Billinghurst. 2008. Trends in augmented reality tracking, interaction and display: A review of ten years of ISMAR. In *2008 7th IEEE/ACM International Symposium on Mixed and Augmented Reality*. IEEE, 193–202.
- [63] Steven Zimmerman, Alistair Thorpe, Chris Fox, and Udo Kruschwitz. 2019. Privacy Nudging in Search: Investigating Potential Impacts. In *Proceedings of the 2019 Conference on Human Information Interaction and Retrieval*. 283–287.