Understanding and Mitigating Technology-Facilitated Privacy Violations in the Physical World

Maximiliane Windl
LMU Munich
Munich, Germany
Munich Center for Machine Learning (MCML)
Munich, Germany
maximiliane.windl@ifi.lmu.de

Verena Winterhalter
LMU Munich
Munich, Germany
verena.winterhalter@gmail.com

Albrecht Schmidt
LMU Munich
Munich, Germany
albrecht.schmidt@lmu.de

Sven Mayer
LMU Munich
Munich, Germany
info@sven-mayer.com

ABSTRACT

We are constantly surrounded by technology that collects and processes sensitive data, paving the way for privacy violations. Yet, current research investigating technology-facilitated privacy violations in the physical world is scattered and focused on specific scenarios or investigates such violations purely from an expert's perspective. Informed through a large-scale online survey, we first construct a scenario taxonomy based on user-experienced privacy violations in the physical world through technology. We then validate our taxonomy and establish mitigation strategies using interviews and co-design sessions with privacy and security experts. In summary, this work contributes (1) a refined scenario taxonomy for technology-facilitated privacy violations in the physical world, (2) an understanding of how privacy violations manifest in the physical world, (3) a decision tree on how to inform users, and (4) a design space to create notices whenever adequate. With this, we contribute a conceptual framework to enable a privacy-preserving technology-connected world.

CCS CONCEPTS

• **Security and privacy → Human and societal aspects of security and privacy**; • **Human-centered computing → Human computer interaction (HCI)**.

KEYWORDS

privacy, privacy policies, smart environments

ACM Reference Format:
Maximiliane Windl, Verena Winterhalter, Albrecht Schmidt, and Sven Mayer. 2023. Understanding and Mitigating Technology-Facilitated Privacy Violations in the Physical World. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23), April 23–28, 2023, Hamburg, Germany*. ACM, New York, NY, USA, 16 pages. https://doi.org/10.1145/3544548.3580909

CHI '23, April 23–28, 2023, Hamburg, Germany

ACM ISBN 978-1-4503-9421-5/23/04…$15.00
https://doi.org/10.1145/3544548.3580909

1 INTRODUCTION

Usable privacy research has long focused on improving privacy communication in the online world as it recognizes the need to inform users when sensitive data gets stored and processed. However, technology has long moved into the physical world: Walking on the street and getting recorded by CCTV cameras, paying at the supermarket with a debit card, talking to a friend next to a smart speaker, or merely sitting beside someone who uses their smartphone – all these situations are commonplace but involve technology with sensors and capabilities to collect and process sensitive information. Hence, it should be common practice to inform users about possible privacy violations to hand them back control over their personal data. However, we currently lack an encompassing understanding of which technology-facilitated privacy violations users experience in the physical world. Consequently, we require establishing methods to provide users with privacy information in these diverse contexts.

Research investigating privacy in the Internet of Things (IoT) mostly refers to the IoT as a whole while not naming concrete scenarios and situations [14, 19, 65] or inquires users about a predefined set of scenarios and parameters selected by researchers [13, 22, 39]. Yet, there is scattered research on concrete privacy violations in the physical world through technology, such as shoulder surfing [17, 53], public CCTV cameras [11, 57], or smart home devices [2, 38]. So far, only Chow et al. [13] tried to create a more holistic description by envisioning five contextual parameters for privacy-relevant scenarios in the IoT. However, they are neither validated nor can they cover an encompassing space of privacy-relevant interactions in the physical world with technology. Since there currently is no encompassing understanding, the efforts to create privacy notices in the physical world lack substance. Feng et al. [19] constructed a design space limited to the privacy choices in the IoT. Others elicited possible mechanisms to communicate privacy for IoT devices in the smart home [70], created a Privacy Nutrition Label to support users during the purchasing stage [18], or compared different prototypes of concrete mechanisms [61].

However, we still require a general design space to create effective mechanisms for privacy-relevant interactions with technology in the physical world to overcome these limitations.

This work creates effective mechanisms using an online survey and expert interviews. First, to understand which situations require privacy notices, we conducted an online survey (N=100) in which we probed participants for privacy violations in the physical world through technology. Based on the gained insights, we constructed an encompassing scenario taxonomy for user-experienced physical-world privacy violations. Second, we conducted interviews (N=10) with privacy and security experts from industry and academia, in which we asked them to design privacy notices using the new scenario taxonomy. Thus, we validated our proposed taxonomy and established mitigation approaches for the most relevant scenarios. Based on these insights, we refined our proposed taxonomy and, through analyzing the experts' designs, provide guidance for mitigating technology-facilitated privacy violations in the physical world. In detail, we gained an understanding of how privacy violations manifest in the physical world, composed a decision tree on how to inform the user, and lastly, extracted a design space to develop notices.

Our contribution is fourfold. To the best of our knowledge, this paper is the first to construct an encompassing and validated (1) scenario taxonomy for user-experienced privacy violations in the physical world through technology. This taxonomy will help designers understand which situations require close consideration of privacy mechanisms. Moreover, to establish effective mitigation approaches, we contribute three additional conceptual tools: (2) the dimensions of privacy violations, (3) a decision tree to decide on the best privacy-preserving mechanism, and (4) a design space to develop notices. As such, we contribute a conceptual framework allowing designers and developers to address users' privacy concerns in the physical world.

2 RELATED WORK

We first describe prior research on privacy violations and point out the gap regarding technology-facilitated privacy violations in the physical world as they are experienced from a user's perspective. We then reflect on how these privacy violations are currently mitigated and explain why we need more research on mitigating privacy violations in the physical world.

2.1 Understanding Privacy Violations

An extensive body of research focuses on privacy violations in the *online* world. This includes research on privacy violations in social networks [28, 30, 31], through online photo sharing [26, 27], and privacy violations caused by smartphone sensors [10], notifications [64], and apps [51].

Research investigating privacy violations outside the *online* world is focused on multiple specific violations. One area of interest is public filming and surveillance through CCTV cameras [11, 57] or drones [63, 72]. Also, violations related to shoulder surfing, i.e., someone gaining illicit access to sensitive information by observing a private display, have been thoroughly investigated, e.g., [17, 53]. Due to their placement in the most intimate spaces, privacy violations related to smart home devices have also experienced significant uptake. Such research includes investigating the privacy expectations of smart home visitors [36] and bystanders [71], as well as research on violations caused by listening devices, such as smart home assistants [1, 54]. Specifically, Barbosa et al. [3] found that users were most uncomfortable with information flows that allow conclusions about demographics or that monitor communications or lifestyles. Moreover, Lafontaine et al. [33] found that users were primarily concerned about smart home devices transmitting data without consent or having security loopholes. Lau et al. [34] found that privacy concerns are one of the most significant reasons for not using smart speakers. In this regard, users reported being mainly concerned about the devices always listening and that data might be used for targeted advertising or shared with third parties. In addition, a survey by Windl and Mayer [68] showed that users' concerns vary by device and sensor type. While users are most concerned about microphones and cameras, they have few concerns regarding motion or temperature sensors.

While most research investigating privacy violations in the physical world is scattered across multiple specific areas, there is research employing a more holistic approach. Oates et al. [42], for example, investigated the difference between laypersons' and experts' mental models of privacy by analyzing drawings to the prompt *"What does privacy mean to you?"* Even though the question was framed quite openly, people equally often drew concepts related to physical privacy as they drew concepts related to technology-facilitated privacy, surfacing the significance of protecting users' privacy across both areas. Gerber et al. [23] conducted interviews to investigate users' mental models of privacy consequences and their obstacles and strategies for privacy protection. They found that most users are unaware of possible consequences and refrain from protecting their privacy as they find it too complicated or lack the necessary knowledge. Naeini et al. [39] conducted a large-scale online vignette study to learn about peoples' privacy expectations and preferences in the IoT. They found that privacy concerns vary depending on the context, such as the type of data collected or the location of the collection. For this, they presented their participants with 14 pre-defined vignettes and parameters that the researchers hypothesized would influence peoples' privacy preferences. While this approach allowed the researchers to quantify their results, it restricted the users' responses to the selected scenarios. Similarly, Gerber et al. [22] did a large-scale online survey investigating how users' estimation of the probability and severity of nine risk scenarios differed depending on how concretely the scenarios were described. They found that users underestimate the privacy risks posed by abstractly described scenarios. However, the users' responses were yet again restricted to the scenarios selected by the experts. Furthermore, Chow et al. [13] envisioned a set of five contextual parameters to describe privacy-relevant interactions in the IoT: *Where* the data is collected, *what* data is collected, *who* collects the data, the *reason* for the collection, and the *persistence* of the data collection. They built scenarios using these parameters to ask participants in interviews for their perceptions. They found that mainly the purpose of the tracking and the entity collecting the data influenced participants' perceptions. However, those parameters were again envisioned by researchers and are not based on

experienced privacy violations. As such, they are neither validated nor encompassing. Moreover, Solove [58] constructed a comprehensive taxonomy of privacy harms to understand which activities invade people's privacy and ultimately evaluate the effectiveness of privacy protection measures. Solove [58] describes the problem space along the four axes of information collection, information processing, information dissemination, and invasion. While Solove [58] provides a useful taxonomy helping us gain a detailed understanding of privacy-violating activities, it was likewise envisioned solely from an expert's point of view.

The described works are valuable to understand and especially quantify insights related to users' perception of privacy violations. However, all prior investigations either did not explicitly focus on privacy-violating scenarios users experience in the physical world or were restricted to the parameters and scenarios selected by researchers. This means they do not necessarily reflect actual privacy violations experienced by users, nor can they cover the diverse landscape of technology-facilitated privacy violations. Consequently, we address this research gap with our first research question ***(RQ1)**: In which situations do users feel privacy violated in the physical world through technology?*

2.2 Mitigating Privacy Violations

Already in 2008, Spiekermann and Cranor [59] proposed a framework that distinguishes two general approaches to mitigate privacy violations: *privacy-by-policy* and *privacy-by-architecture*. While *privacy-by-architecture* aims to design a system in a way that does not harm users' privacy in the first place, e.g., by minimizing the collection of sensitive information or anonymizing data, *privacy-by-policy* aims to protect users' privacy by implementing the principles of notice and choice. While they conclude that *privacy-by-architecture* should generally be preferred as it provides a higher level of privacy, many businesses choose to implement *privacy-by-policy* as it does not interfere with their business model of collecting extensive amounts of user data.

Research and lawmakers have long recognized the need to inform users about data practices when sensitive information gets stored and processed in the *online* world. As a result, legal requirements, such as the General Data Protection Regulation (GDPR), EU's ePrivacy Directive (EPD), and the California Consumer Privacy Act (CCPA) exist to prescribe how users must be informed about privacy regulations [46–48]. Privacy policies and cookie banners are two current efforts to meet these requirements. However, prior research found that privacy policies have several hurdles, such as their length [37, 43], difficult legal language [35], and abstract wording [50], which leads to people finding them overwhelming [45] and ultimately making privacy policies insufficient in providing users with notice and consent. Therefore, the community focused on improving them either through developing standards [52] and guidelines [49, 55] or supportive tools and visualizations [32, 62, 66]. One example deployed in practice is a visualization developed by Kelley et al. [32] that adopts the principles from nutrition labels for privacy policies. Apple and Google meanwhile require all apps in their app stores to have such Privacy Nutrition Labels[1], even though these labels recently experienced criticism as they are not prominently placed, use confusing terminology, and are inconsistent with the apps' privacy policies [15]. Another promising example are Contextual Privacy Policies that, along with the principle of Contextual Integrity [40], provide only the relevant snippets from the privacy policy in the context where they are relevant [66]. In contrast to privacy policies that mostly have to be actively retrieved, cookie banners are automatically displayed and require the users' active consent. However, most use dark patterns to nudge users towards consenting to the less privacy-preserving options, such as hiding privacy-controls at the end or on a sub-page [24, 25, 41]. Here, Nouwens et al. [41] found that if navigating to a control option took effort, they were seldom used. On the other hand, if granular control options were available right away, user consent decreased by 8-20 percentage points. As such, communicating privacy *online* has been thoroughly investigated, and protecting users is even required by law.

Common modern personal computing devices, e.g., smartphones, laptops, and computers, are equipped with sensors that can collect sensitive data, making protecting users' privacy even more important. Although manufacturers recognized this need and implemented control options, research found that most users do not know where to find them or how to use the settings effectively to protect their privacy [6, 16, 20]. Thus, recent research is actively investigating how to improve the communication of privacy information, for example, by developing supplementary apps or integrating privacy dashboards [4, 21]. As long as users stay on websites or apps, the privacy concerns are limited to the *online* world. However, the boundaries between *online* and *physical* worlds can easily overlap, for example, when the phone is used to take pictures in public or when a voice assistant recognizes voices.

In regards to privacy in the *physical* world, research has paid special attention to smart homes. Here, data gets collected in the most intimate spaces and, thus, is subject to even higher expectations [3, 34, 44]. Through a co-design study, Yao et al. [70] investigated how users design privacy mechanisms in smart homes and found that they relied on rather simple strategies, such as disconnecting from the internet or introducing a private mode. Thakkar et al. [61] developed and compared four different visualizations, including ambient smart lights and a privacy dashboard. They found that different visualization have different pros and cons and, thus, are suitable in different contexts. While the ambient light, for example, provided unobtrusive information, the data dashboard enabled detailed insights. However, smart devices are not only present in homes but can meanwhile be found in nearly all areas of daily life. Personalized privacy assistants have been recommended as an effective means to communicate privacy information in the IoT [7]. They are envisioned to learn the users' privacy preferences, adjust settings automatically, and make privacy-relevant decisions on the user's behalf. Users were generally positive about personalized privacy assistants but also expressed concerns regarding balancing having more awareness and control while not being overwhelmed by notifications [14]. To create a more holistic understanding, Feng et al. [19] constructed a comprehensive design space for meaningful privacy choices in the IoT with five dimensions: type, functionality, timing, channel, and modality. However, this space is limited to privacy choices.

[1] https://www.apple.com/privacy/labels/, https://blog.google/products/google-play/data-safety/

In sum, research has intensively investigated how privacy can be protected *online*, as preserving users' privacy is even prescribed by law. However, in regards to mitigating privacy violations in the *physical* world, current research is scattered and focused on specific use cases, such as smart homes. Consequently, we still lack an encompassing understanding of how to protect users' privacy in a multitude of other scenarios in the physical world. Therefore, we ask the research question ***(RQ2)**: How can we effectively mitigate privacy violations in the physical world?*

3 UNDERSTANDING USERS' PHYSICAL-WORLD PRIVACY NEEDS

Prior work investigated users' privacy concerns using different scenarios and parameters. However, these factors were all constructed by researchers, and the investigations were limited to the selected scenarios, see Section 2.1. In contrast, we aim to construct a more extensive taxonomy of technology-facilitated privacy violations in the physical world, as they are experienced **from a user's perspective**. With this, we want to understand which scenarios users perceive as privacy-relevant and, thus, require effective mitigation approaches. We conducted an online survey on Prolific to answer our first research question (**RQ1**). We provide all participant responses to fellow researchers upon request.

3.1 Survey Construction

In our survey, we asked participants to provide at least three scenarios but as many as they could. We asked them to envision possible new scenarios if they could not remember three they had experienced. See Appendix A for the complete questionnaire.

We first created an initial draft with a set of questions that we piloted with colleagues (N=6). As we did not want to bias our participants to get a most encompassing set of diverse scenarios, we did not provide concrete examples of privacy violations in the physical world. However, without concrete examples, it was hard for our participants to grasp the meaning of privacy violations in the physical world. Therefore, we added a short introductory paragraph describing what privacy violations in the physical world mean – as concretely as necessary but as vague as possible. Next, we did a second round of testing with members of our university (N=42). Now, the main problem was that participants frequently provided scenarios that referred to the usage of web pages or apps or were very vague in their responses. Therefore, we added several times throughout the survey that we do not refer to privacy violations caused by web pages or apps and added more detail to our question. Ultimately, the final questionnaire consisted of three blocks: 1) an informed consent form, 2) demographic questions, and 3) the main questions of the survey. After we piloted this final set of questions on Prolific (N=24), we received sufficiently concrete responses. The wording of the main questions is as follows:

> *Have you ever felt like your privacy was violated in the real world by means of technology?*

If participants indicated that they had never felt that way before, we asked them to envision scenarios where technology could violate their privacy in the real world. We worded the question as follows:

> [If they *had* experienced privacy violations before]
> *For the situations in which you have felt that your privacy was violated by technology, describe the scenarios as concretely as possible, including: 1) where you were when the violation happened, 2) which technology caused the violation, and 3) what kind of private information was affected.*
> [If they had *never* experienced privacy violations before]
> *Envision real-world scenarios where your privacy can be violated by technology. Describe the scenarios as concretely as possible, including: 1) where you are when the violation happens, 2) which technology causes the violation, and 3) what kind of private information is affected.*

3.2 Participants

For the main survey, we recruited 100 participants on Prolific. We excluded 14 during the analysis of the statements since we agreed after a thorough discussion with three authors that their provided scenarios did not sufficiently describe privacy violations in the physical world. Of the remaining participants, 46 were male and 39 female with an age range from 18 to 57 (M = 34.8, SD = 9.5). The participants resided on five continents: Europe, Africa, Asia, South America, and North America. Specifically, the five most represented countries were Spain (12), Ireland (9), the United Kingdom (8), Greece (7), and Portugal (7). All participants were full-time employees with various professions. The five most represented sectors were IT (11), hospitality (6), education (5), law (4), and retail (4). The participants spent, on average, 13 minutes (SD = 8.7min) to complete the survey and were compensated with 1.90£

.

3.3 Data Analysis

In total, we received 268 scenarios (on average 3.1 scenarios per participant, min = 3, max = 8). Of those, 153 were experienced, and 115 were envisioned privacy violations. We used thematic analysis and Atlas.ti to make sense of the data [5]. First, three researchers independently open-coded a random subset of 20% of all scenarios. Afterward, we discussed our initial codes in person to resolve ambiguities and create an initial code book. Next, we divided the remaining statements among us for coding. Finally, we met again to discuss our codes and formed code groups and overarching themes. We repeatedly discussed and reformed those themes by comparing the coded snippets with coded extracts from the other code groups. During this process, we excluded 75 scenarios as the violations were either not caused by technology (e.g., someone seeing one's address on an envelope), described privacy violations caused by using web pages or apps (e.g., targeted ads after doing Google searches), or described actual frauds or scams (e.g., cloning credit cards at ATMs). The following are examples of included scenarios provided by our participants:

> *"Because of a shopping center fidelity card, the company sent a list of bonuses to use in future visits to the shops, showing that they know exactly what I used to buy, when, how much, and new things I started to buy. I felt*

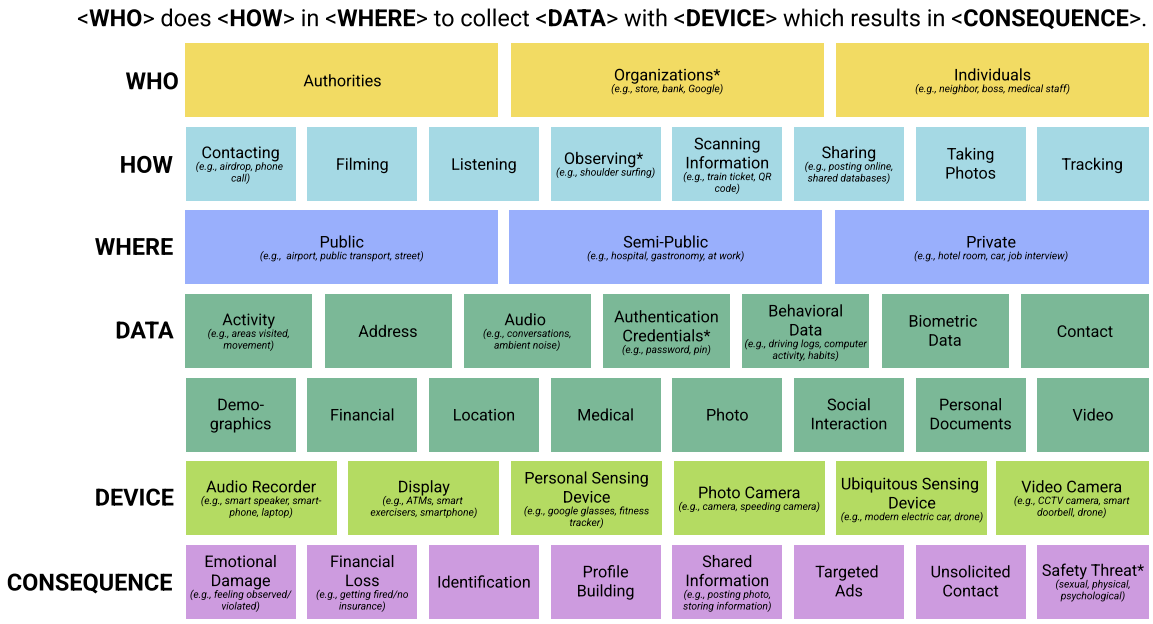

Figure 1: The scenario taxonomy. Scenarios can be built using the sentence structure "<WHO> does <HOW> in <WHERE> to collect <DATA> with <DEVICE> which results in <CONSEQUENCE>" and filling the placeholders with one or multiple of the individual building blocks. Note: The building blocks marked with a * were carefully reformulated in response to the expert interviews but do not distort the users' perspective.

> *they knew exactly how I lived and did things in my private life."* — P15
>
> *"I felt my privacy was violated when a neighbor of my old house put security cameras around his place, but they were recording my house too, recording every movement of mine and my family's."* — P42

First, we created 317 unique codes. Then, after multiple iterations in hour-long sessions, we settled on a final set of 7 themes with 43 code groups and 188 unique codes.

3.4 Awareness and Voluntariness

AWARENESS AND VOLUNTARINESS is an overarching theme that spans all scenarios, giving reflections on a meta-level. Namely, whether participants were aware of a violation happening and whether they voluntarily exposed themselves to situations. In regards to being *forced*, multiple participants mentioned how their information was acquired without their consent: *"[...] They took a photo of me without my consent [...]"* (P6). In contrast, participants also recited occasions where they *voluntarily* exposed themselves to situations. Here, P31, for example, recalls a situation where they did not mind when their boss installed CCTV cameras in the changing rooms: *"[...] I had no issue getting changed there, but other people might have."* However, to provide data and information voluntarily, the person has to be *aware* of the information being gathered in the first place. Here, several participants recall situations where they were completely unaware and only learned afterward that they had been exposed to a privacy violation. For example, P94 remembers a situation where they did not know how a digital doorbell recorded their audio and video: *"Afterwards, I heard that I was being recorded in both audio and video while standing in front of the door. There was however no light indicating this."*

3.5 Scenario Taxonomy for Physical-World Privacy Notices

The remaining six themes describe the scenarios in detail. Hence, we used them to construct a scenario taxonomy for user-experienced physical-world privacy violations, see Figure 1. Those six themes were: WHO, HOW, WHERE, DATA, DEVICE, and CONSEQUENCE. The following sentence structure can be used to build concrete scenarios using our taxonomy:

> <WHO> does <HOW> in <WHERE> to collect <DATA> with <DEVICE> which results in <CONSEQUENCE>.

Each theme has multiple building blocks that describe its individual dimensions, see Figure 1. It is important to note that these building blocks are not necessarily exclusive. A privacy violation can, for example, result in a user's *identification*, which can additionally cause *emotional damage* as the user might feel observed or spied on. In the supplementary material, we provide a list of all our respondents' examples for each building block. In the following, we describe the themes in more detail.

WHO illustrates the entity collecting the data. It consists of the three groups *government*, *organizations*, and *individuals*. *Government* encompasses all legal entities. As such, it differs from the other two groups since its actions are subject to official regulations. P70, for example, provided a scenario where they felt violated through a governmental institution when crossing a border: *"I know that that data will be stored and it may be used for whatever purpose by*

the US government." *Organizations* encompass all entities where the individual can not be identified. This could be a store, an electricity supplier, or a big tech firm like Google. Here, P70, for example, reported feeling violated by *"Google Maps cars that map the street"*. *Individuals* enclose all entities, where we can identify a single person as the data collector. Even though we cumulate *individuals* in one group, this group is quite diverse, containing various social relationships, ranging from delivery drivers, bosses, and neighbors to friends. For example, P74 felt violated when their identity card was scanned by *"delivery drivers asking for proof of identity."*

How describes the way in which the data is collected. It contains eight different modes of data collection, including *taking photos*, *filming*, *scanning*, *sharing information*, and *observing*. P87, for example, provided a scenario where they felt privacy violated through their license plate being scanned: *"When entering residential estates in South Africa [...] my vehicle license disc [is] scanned by the security company at the gate [...]."* In contrast, a teacher felt violated when their *"students took pictures [of her] during breaks at school (P12)."*

Where describes the space in which the violation happened. It can be *public*, *semi-public*, or *private*. *Public* encompasses all the spaces everyone can enter without access restrictions. P70, for example, felt violated while walking in a city: *"In Seoul, there were a lot of CCTV cameras [...]. And sometimes I felt "observed"*. *Semi-public* are all places with some access restrictions, such as requiring a ticket or special permission, but which are also not private since they are shared with multiple people. Such a semi-public place can, for example, be a backyard where P57 reported feeling violated by *"home security cameras filming the entire backyards for security purposes."* *Private*, in contrast, are spaces where the person is usually alone or with familiar or trusted people, such as a hotel room or a doctor's examination room. For example, P40 felt violated in their home as a *"security system was recording what [they were] doing."*

Data illustrates the kinds of data collected, whereby we distinguish between 13 groups of different data types. Some of the data types are related or directly connected. For example, location, activity, or behavioral data can be derived from a video stream. P16, for example, is concerned about the various information that can be derived from a speeding camera's footage: *"How fast I am traveling, in what hour, in what vehicle, and where."*

Device represents the entity collecting the data, whereby we distinguish the device groups by their main sensor or capability evoking the privacy concern. Therefore, we also have the group *display* since participants mentioned scenarios where their personal information was exposed on a publicly visible screen, such as their account balance displayed on an ATM, as P28 reported: *"I used to hate it when ATMs used to show my balance and full name [...]."* Personal sensing devices encompass all devices worn by a person that mainly gather their users' data. In contrast, ubiquitous sensing devices are placed in the surroundings. P43 provided an example of a ubiquitous sensing device, in the form of a modern electric car that sends driving logs to the car's manufacturer.

Consequence describes the result of the privacy violation. These can be concrete results, such as *targeted ads* or *financial loss* or more subjective experiences, such as experiencing *emotional damage*, for example, feeling concerned, unsafe, or violated. P82, for example, reported a situation where they experienced emotional damage through a camera in their workplace's canteen: *"[...] it was uncomfortable knowing that even on our break times our movements (if not our conversations also) were being recorded."*

4 TAXONOMY EXPLORATION: EXPERT INTERVIEWS

To validate and explore our scenario taxonomy and establish concrete mitigation strategies (**RQ2**), we conducted semi-structured interviews, followed by a design session, with ten domain experts from industry and academia. We chose to conduct in-depth interviews instead of a survey to gain a deep understanding of the dimensions of privacy violations, which is not possible with only numbers or short textual responses – the kinds of data typically collected from surveys.

4.1 Interview Protocol

We asked all participants to fill out a short survey to gather demographics and background information before the interview. Before we started the audio recording, we ensured that the participant had filled out the pre-interview survey and the informed consent form we sent via email. The interview was structured into two parts.

In the first part, we explored our taxonomy by letting the experts create scenarios and asking for feedback afterward. For that, we first defined privacy violations in the physical world through technology by giving a concrete example and explaining that we exclude all violations caused by web pages or apps. We then showed and explained the initial version of the scenario taxonomy, after which we asked for any initial feedback. We tasked our experts with creating four distinct scenarios using our taxonomy that they considered especially interesting and privacy relevant. While they created the scenarios, we asked them to use the think-aloud protocol [56]. After, we asked for each scenario to which extent they agreed with the following statement: *I think this scenario is very concerning* on a Likert scale from 1 to 7. Finally, we discussed the scenarios in more depth along the dimensions of voluntariness and awareness, as we had discovered those as important axes during the analysis of the survey statements (see Section 3.4).

The second part of the interview was an interactive design session where we tasked participants with creating a privacy-enhancing mechanism/solution for their scenarios using their preferred sketching tool. We again used the think-aloud protocol [56] as they created the mechanisms. After they had finished a design, we asked them to explain in depth how the mechanism works, including how much information is given to the users, when they are informed, how they get notified, which modalities are used, and so on. Finally, we wanted our participants to think about how their mechanisms influence the impact of the privacy violation. For that, we asked them to re-evaluate the statement of *I think this scenario is very concerning*, imagining their mechanisms were in place and to think about how their mechanisms shift the violation regarding voluntariness and user awareness of the data collection.

In the end, we asked for any additional feedback, thanked our participants, and acquired contact details for compensation. We provide the complete interview guideline in the supplementary material. An interview took approximately one hour, and we compensated our experts with 10€ per hour.

Table 1: Our experts' demographics: Their countries of residency, their continent of birth, their gender, age, highest educational degree, years of experience in their current role, and whether they work in academia or industry.

PID	Residency	Birth	Education	Exp.	Sector
1	Germany	Europe	Doctoral	3 yrs	Academia
2	Germany	Africa	Doctoral	3 yrs	Academia
3	Germany	Africa	Doctoral	3 yrs	Academia
4	USA	South America	Doctoral	1 yr	Industry
5	Germany	Europe	Doctoral	7 yrs	Academia
6	Germany	Europe	Doctoral	9 yrs	Academia
7	Germany	Europe	Master's	1 yr	Academia
8	USA	Asia	Doctoral	13 yrs	Academia
9	UK	Asia	Master's	2 yrs	Academia
10	Germany	Europe	Master's	3 yrs	Academia

4.2 Participants

We recruited the experts using convenience sampling, followed by snowball sampling. To qualify as an expert, the participants had to match our inclusion criteria which were 1) they must self-identify as a privacy expert, and 2) they must currently pursue or have pursued a Ph.D. in a privacy-related field.

The ten privacy experts who took part in the interviews were aged between 28 and 45 years ($M = 32.9$, $SD = 4.9$). Most (9) experts worked in academia while one was working in the industry. The experts had different cultural backgrounds. Most (5) were from Europe, two from Africa and Asia, and one from South America and India, respectively[2]. The experience levels of our privacy experts in their current role ranged from a minimum of 1 to a maximum of 13 years ($M = 4.5$, $SD = 3.9$). For a detailed overview of our experts' demographics, see Table 1.

4.3 Data Analysis

In total, we recorded 10.73 hours of audio data (on average 64.4 *min* per expert, $SD = 10.1 min$) and used Atlas.ti and thematic analysis to analyze our data [5]. For that, two researchers first independently open-coded two randomly selected interviews, after which the researchers met to discuss the initial codes and form a joint code book. The remaining interviews were then divided amongst the researchers for coding. A third researcher joined the group for the third iteration, where we formed groups of related codes and overarching themes. We repeatedly reworked and refined these themes by comparing the coded snippets across all themes. To analyze the privacy designs, we synthesized all designs and extracted the design factors together with our experts' remarks about their designs. Through multiple iterations and discussions, this process led to 4 themes with 21 code groups and 242 unique codes. Those four themes were: Insights on Scenario Taxonomy, Dimensions of Privacy Violations, Practicality of Privacy Notices, and Design Factors.

[2]We grouped the country of birth by continent to ensure our participants' anonymity.

5 FINDINGS OF THE EXPERT INTERVIEWS

We present the results from our expert interviews along our four themes. Insights on Scenario Taxonomy encompasses all our experts' feedback on the taxonomy, especially regarding missing or misleading blocks and ambiguous terms, which we used to update the taxonomy. Dimensions of Privacy Violations describes our experts' reflections on the notions of voluntariness and awareness, for example, what constitutes a violation and the different levels of being aware. In Practicality of Privacy Notices, our experts reflected on when privacy by design should be preferred over giving a notice, which we used to create a decision tree to decide on the most suitable mechanism. In Design Factors, we synthesize all our experts' notions about their designs with the factors we extracted from their sketches to construct a design space for privacy notices in the physical world.

5.1 Insights on Scenario Taxonomy

Most experts were excited about the taxonomy and found it encompassing and sufficiently descriptive (E1, E3, E5, E7, E9). E9, for example, stated: *"The taxonomy looks really nice, and I like how you have presented it. I think it conveys the complete picture."* In addition, all experts successfully created 3 or 4 scenarios using our taxonomy, validating its applicability. However, the scenario-creation process also revealed one missing building block (Authentication Credentials, e.g., passwords, pins) and opportunities to improve wording, as sometimes experts had to ask to clarify what they described. Therefore, we reformulated four building blocks and added one new block to allow for the overall better usability of this conceptual tool. As we did not want to dilute the users' perspective nor present the experts' perspective, we marked all blocks, which we reformulated in response to the expert interviews with a * in the taxonomy, see Figure 1. We detail all reformulations below.

Organizations. We renamed the block from *companies* to *organizations* as several experts had trouble fitting their examples in. E5, for example, mentioned that they would not know where to put NGOs as they are neither an authority nor a company. The same applies to political parties as they are technically not a company. As a result, we renamed this entity to organizations to also account for other groups of individuals.

Observing. Another adjustment we made was that we renamed the block *shoulder surfing* to *observing*. This was the response to E4 creating a scenario that initially did not fit our taxonomy: *"Border agent obtains your device in an airport to collect and check your conversations and emails in order to check for suspicious activity [...]."* While this is not a hidden act, as is *shoulder surfing*, the violation is also caused by viewing the screen. Hence, we renamed the block to *observing* and made *shoulder surfing* an example of this block.

Authentication Credentials. This was a new block we created, as E6 pointed it out as missing. They remarked that they would not know where data *"like security, critical data, like authentication credentials, a PIN, and so on"* would fit. As this was not an example that came up in our survey but is data that could be exposed through privacy violations in the physical world, we added it to our taxonomy.

Safety Threat. The last block we added was safety threat. Two experts brought this up as they constructed scenarios with far more

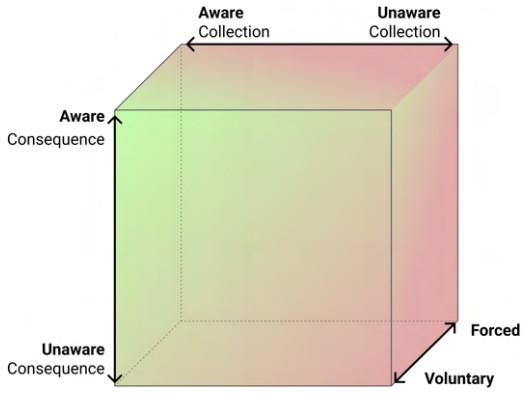

Figure 2: The dimensions of privacy violations in the physical world. All violations in the physical world can be placed along the dimensions. There are three continua: voluntary ↔ forced, aware of data collection ↔ unaware of data collection, and aware of consequence ↔ unaware of consequence. The color coding depicts a state's desirability, whereby green means desirable and red means undesirable.

outreaching consequences than *emotional damage*. For example, E1 constructed a scenario where a user's presence at home was deducted from a smart thermostat and then exploited to break into their home, and E8 created a scenario where technology led to stalking in the physical world – a far more profound consequence than *emotional damage*. Thus, we added *safety threat* to our taxonomy.

Another more general feedback from our experts was that some of the building blocks could be ambiguous and contain different nuances (E1, E8, E10). E8, for example, mentioned trust as an important layer in the *who* category. While E8 generally liked the three groups, they remarked that elements of a building block could be judged differently depending on the level of trust:

> *"I think the who at a high level... these are the entities, but within them, I think, there are definitely varying familiarities, trust, previous experience. They can be subdivided."*—E8

The notion of ambiguity also came up regarding *where*. E8 and E10 noted that whether a space is considered private or semi-public can vary and depend on individual judgment:

> *"For private, I think it's a huge difference if it's my own private space or if it's the Airbnb example, which is, I'm not sure if you consider that a private space as well."*—E10

This ambiguity of some building blocks should be considered when evaluating scenarios as they can influence how violative people perceive a situation. Yet, as these factors do not affect the ability to form scenarios, we did not incorporate them in our taxonomy.

5.2 Dimensions of Privacy Violations

When discussing the dimensions of privacy violations, E4 reflected on when a privacy violation should be considered a violation. In this regard, E4 discussed that whenever people get into a situation completely voluntarily and aware of what is about to happen, it should not be considered a violation. Furthermore, experts discussed how awareness again has two dimensions. First, a person can be aware of the data being collected. Second, a person can also be aware of the consequences, i.e., what it means that a certain data type gets collected (E4, E8). E8 explains this in detail:

> *"They may have just resigned to it, saying, 'Okay, I kind of voluntarily give it because I want the benefit of, let's say, social media or social interactions.' What they're not seeing is the potential consequences and harms because the probability of those harms materializing could be small, or the consequences might happen six months later, and they want the benefit now, or it may be that they're looking at the benefit of this one action, and they don't necessarily understand that hundreds of these actions put together do reveal a lot more than simply that one action."*

Finally, several experts discussed how the dimensions of forced and voluntary and aware and unaware should not be seen as binary but more as a spectrum as they also contain several levels (E1, E2, E3, E4, E8). As a concrete example, E4 discussed a situation where someone placed cameras in a changing room. While the person might be aware that there are cameras, they might not know where exactly they are placed, creating a state of in-between awareness. However, this is not only the case for being aware of the data collection but also about the consequence. While a user might have some knowledge that collected data gets stored and transmitted somewhere, they might still not know what that means for the bigger picture, as E8 explains:

> *"You can totally be aware of what happens after the collection in terms of, oh, it goes to the cloud, and it's stored there, and it's retained for whatever, and, I don't know, three people have access to it, but you still might not know the consequences of what can happen from that."*

Additionally, there are also stages between voluntary and forced, as E4 pointed out:

> *"There are things that may not necessarily be completely forced. It might not be, 'I have to disclose my vaccination status to maintain my employment,' but it could be, 'If I don't disclose my vaccination status, I am not allowed to work in the office.' I have an option to do it. It's not just like yes or no."*

All these notions and insights helped us create the dimensions of privacy violations in the physical world as depicted in Figure 2. The figure has three axes as identified by our experts: (1) being aware of the data collection happening, (2) being aware of the consequences, and (3) engaging in a situation voluntarily. All these axes are continua, as depicted by the double arrows. The colors depict the desirability of a state, whereas green means desirable and red means undesirable. Hence, being aware of the data collection, the consequences, and the opportunity to engage in a situation voluntarily is the perfect state for the user. It is important to note that the upper right edge, i.e., a state of being unaware of the collection but aware of the consequence, does not exist, as reflected in the figure by the missing border.

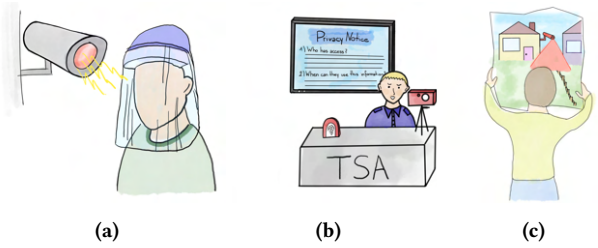

Figure 3: The re-sketched designs of our experts: a) shows E8's design of a hat with a face shield protecting its wearer from face identification; b) shows a display as designed by E1 to inform during the immigration procedure; c) shows a paper plan as designed by E4 to visualize which parts of one's property are captured by a neighbor's CCTV camera.

5.3 Practicality of Privacy Notices

During the second part of the interview, we asked our experts to sketch out notices for the scenarios they had constructed using our taxonomy. While our experts created the sketches, they reflected on the practicability of privacy notices. Several experts did not consider a notice the best approach as it gives the burden to the user (E4, E6, E8). Instead, they suggested focusing on preserving users' privacy by default by not collecting or obscuring the collected data, as E6 explains:

> *"Instead of notices, in these situations, I would try focusing on solutions that don't take the data in the first place. Let's say we have the video situation, and there are ways of anonymizing videos and having the source encrypted somewhere. If there is some crime happening, you can say I need the real sources because I need to identify the people. If you just use it for statistics like how many people are walking there and so on, you can use the anonymized version."*

However, as this quote already indicates, obscuring the data is sometimes impracticable. For example, in security-sensitive scenarios, such as when a border agent acquires biometric data. In these cases, when the user is forced to comply, and data obfuscation is impossible, *"raising awareness might be the only option"* (E1), and as such, a privacy notice is the most sensible approach. However, notices can also easily become overwhelming, especially when the data collection is constant and ubiquitous, such as in cases of public surveillance through CCTV cameras (E4, E6, E8, E10). Although a phone could, for example, always vibrate when the user is in the vicinity of a camera, this could quickly result in the user being overwhelmed and unable to deal with the massive amount of notifications. In these cases, notices might even have *"a counter effect in making people nervous while they can not do anything"* (E4) and, thus, should be avoided. The next important question is whether the user can control the technology causing the violation. If this is the case, the user can employ mitigation strategies, such as unplugging or turning the device off. E2, for example, designed a simple plug as a privacy-preserving mechanism. In addition, researchers can support users by developing assistive tools that ease controlling devices, improve the overview, or develop devices that give control to the users. E10 designed such a mechanism in the form of a smartphone application that not only notifies the user upon entering a smart home but also allows users to deactivate them. However, other tools are needed if the user has no control over the technology. We coined these "guerilla tools" as they employ exotic strategies to obscure or destroy data during collection and are actively leveraged by the user. An example of such a tool was brought up by E8, who designed a hat worn by the user that sends out a signal to destroy the images captured by face recognition cameras, see Figure 3a.

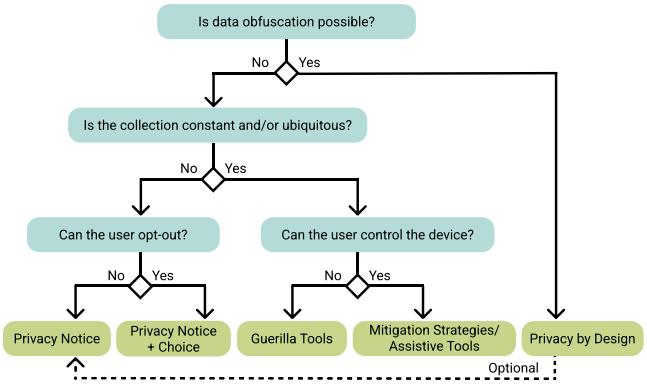

Figure 4: Decision tree to decide on the best privacy-preserving mechanism in the physical world based on the constraints of a situation.

Out of these insights, we created a decision tree for privacy-preserving mechanisms in the physical world, as depicted in Figure 4, leading to the best mechanism from a user's perspective. The tree works as follows: First, we need to decide whether anonymizing data, i.e., employing privacy by design, is possible. This is the first decision, as our privacy experts made clear that privacy by design reflects the optimal case by preventing a situation from being violative in the first place. As such, it lifts the burden of taking care of privacy-preserving measures from the users. However, obscuring and destroying data is sometimes impossible as information needs to be stored and remain un-anonymized, such as biometric data collected while crossing country borders. As in this specific case, the fingerprint does not get deleted in a timely manner, principle 5 of privacy by design according to Cavoukian [9] is violated. Hence, whenever privacy by design is not viable, data will eventually be collected, and as such, users should be informed. Here, the next decision is whether a violation is constantly happening, such as for public surveillance through CCTV cameras[3]. This is important to clarify as constant and ubiquitous data collection would quickly lead to notices becoming overwhelming and, thus, useless. If we decide that notices are not viable, users need to employ privacy-preserving mechanisms themselves. Here, the third decision is whether users have control over the device collecting the data. Since then, they can employ mitigation strategies, such as unplugging a device. Yet, if a user can not control a device, the only option left to mitigate

[3]We note that in this case, a privacy notice, in the form of a sticker or note in the physical environment, represents an additional option. However, as our decision tree leads to the *best* privacy-preserving mechanism, we do not recommend such signs as they are criticized as representing mere warnings instead of real notices [8].

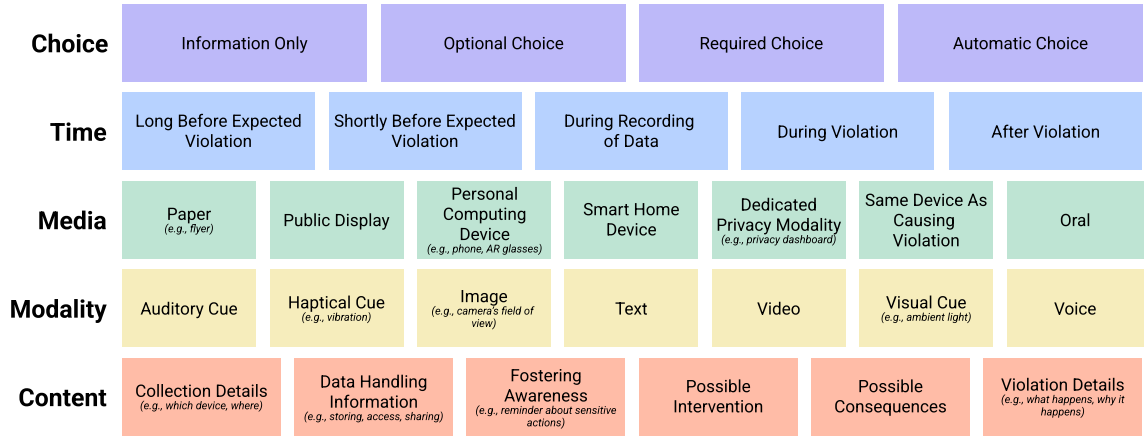

Figure 5: Design space for privacy notices in the physical world.

concerns is to employ guerilla tools that obscure or destroy the data. When we find at the second decision that the data collection is not constant and notices are a suitable approach, the last question is whether the user has a choice; since then, we should provide notice and choice as it hands autonomy over their personal data back to the user.

5.4 Design Space for Privacy Notices

We joined our experts' statements about their designs with the factors we extracted from their sketches. For that, we systematically went over our experts' designs and noted: (1) the general form of notice they used and whether they provided choices, (2) at what time the notice was displayed to the user, (3) the media they used for displaying the notice, (4) the modality to convey the information, and (5) the exact content of the notice. Hence, we extracted five key elements of privacy notices in the physical world: Choice, Time, Media, Modality, and Content. We compiled these factors into a design space, as depicted in Figure 5.

Choice is the first thing to decide when creating a privacy notice as it majorly influences the two following factors time and content. When not providing choices but restricting the notice to informing, the user has no options to control a possible privacy violation. As such, a notice *"might not be the most effective way of protecting privacy, but is still a great awareness tool"* (E4). Consequently, providing choices should be preferred whenever possible. Yet, providing choices is subject to certain constraints as described previously (see Figure 4). Overall, there are three different types of choices. *Optional choices*, which do not force users to interact by having a default option. For example, E8 designed such a choice in the context of an automatic passport scanner, where a screen provided the option to click *no*. However, when the user would not interact with the screen, the default situation of automatic passport control takes effect. *Required choices*, on the other hand, force the user to decide before being able to proceed. The last form is *automatic choices* that decide based on users' previously entered preferences.

Time includes five possibilities on when to deliver a notice. A notice should be given long before a violation if it requires a certain time to react upon. E4, for example, designed a notice for when a border agent checks conversations on a mobile phone. As a solution, E4 designed paper flyers to be handed out while people were still on the plane to give them adequate time to sort through their messages. Giving the notice shortly before the violation might be enough if protecting oneself does not require a lot of time or if control options are given together with the notice, as is the case with *notice and choice*. On the other hand, giving the notice *during* or *after* the violation only makes sense when the user has no options to avoid the violation. E10, for example, designed AR glasses that inform the user about data practices while going through the security scanner. Giving a notice *during data recording* mostly makes sense when people are to be reminded that their actions could lead to a privacy violation. As such, the *notice* can serve as a prevention tool. Here, E3 designed a notification for when children were recognized in the camera's field of view while taking a picture.

Media describes seven ways a notice can be displayed. Several experts discussed how the presentation of the notice should match the users' mental model of what they expect in a specific situation (E1, E4, E8). This resulted in several experts using the same technology that caused the violation to display the notice for at least one of their scenarios (E1, E3, E4, E5, E6, E7, E9) or using another modality that fitted the context. E1, for example, designed a large screen next to a border agent to explain data collection practices during the immigration procedure. They used a screen as people are used to receiving information via screens in airports, see Figure 3b. Other experts designed a dedicated privacy modality. This could be a more conventional feature, such as a light indicating that recording is in progress (E2) or a dedicated privacy device. E5, for example, designed a device devoted to showing all the information collected by the machine used to top up money for the cafeteria.

Modality describes seven different ways to convey privacy information. This can be an auditory, haptic, or visual cue. Since such cues do not contain a lot of information, they were mostly used for notifications, and our experts often combined them with text (E1, E2, E3, E9, E10). Another modality with a way higher information density was, for example, a video designed by E10 to inform about the capabilities of a smart electricity meter. On the

other hand, an image was designed by E4 in the form of a plan to show which part of a garden is in the field of view of a neighbor's CCTV camera, see Figure 3c.

Content can be divided into six different groups. *Collection details* describes all the information around how the data gets acquired, such as which devices are involved, where those devices are placed, and which kinds of data get collected. *Data handling information* clarifies, among others, where the data gets sent to and where it is stored, as well as who has access to the data. *Fostering awareness* includes reminders that an action can lead to a privacy violation and in which cases the collected data can be used. *Possible intervention* includes concrete tips, such as *"cover up the numeric keypad"* (E2) while entering a pin on a machine or in the case of a *notice and choice*, the choices available to the user. *Possible consequences* does include not only negative consequences but also possible benefits. In this regard, E10 mentioned how they would focus on describing the information collected while not being too negative about the consequences if users have no choice. Concretely, her design was a notice for a digital key that collected how long and when someone was in the office. And lastly, *violation details* describes what, why, where, and how a violation happens.

6 DISCUSSION & IMPLICATIONS

In the following, we discuss our findings and show how our tools can be used to create privacy-preserving mechanisms for the physical world.

6.1 It Is Not a Violation if It Is Voluntary and Aware: The Dimensions of Privacy Violations

In our first research question (RQ1), we asked which situations made users feel violated regarding their privacy through technology. We constructed a scenario taxonomy for violations in the physical world to answer this question from our online survey responses. However, our results indicate that the taxonomy alone is insufficient, as the perception of what constitutes a violation is very subjective. Participants reported occasions where they felt especially violated as they were unaware of data being recorded. On the other hand, participants also recited occasions where they voluntarily exposed themselves to situations that would clearly be privacy-violating for others – yet, they were completely comfortable with it. This shows how different people have different comfort levels regarding their privacy; while a situation might be perceived as violative by one person, it might be completely fine for another. The importance of adjusting privacy to individual needs was also already frequently discussed in related work [29, 60, 69]. Our experts echoed this by stating that as long as something is aware and voluntary, it should not be viewed as a violation. This was also already stated by Solove [58]: *"If a person consents to most of these activities, there is no privacy violation."* However, he also raises the question of what constitutes valid consent. We tackle this with the third dimension of being aware of the consequences in Figure 2: Only when being aware of the consequences can a user give valid consent. Considering this, *our dimensions can help decide when a notice is necessary and when it can be omitted.*

The two notions of awareness and voluntariness are interdependent. For something to happen voluntarily, the user must be aware of the data collection and its consequences. As such, *an effective privacy notice can shift a violation from unaware to aware.* Yet, to shift a violation from forced to voluntary, the user must also have a choice, and that choice must be respected. Consequently, *only when combined with choice can a notice shift a violation to the desirable state of "aware and voluntary."*

6.2 A Notice Might Not Be the Best Approach: The Tradeoff Between Distressing and Providing Awareness

Our second research question (RQ2) explored how privacy violations in the physical world can be mitigated. While designing the privacy pre-serving mechanisms, several experts discussed how they did not think a notice was the best approach. Instead, they explained how privacy by design [9] should be preferred whenever possible as it lifts the burden from the users and prevents a situation from being violative in the first place. However, there are also situations where data obfuscation is not an option as the data needs to remain un-anonymized. But even then, notices are not necessarily the best solution. As soon as the collection is constant and ubiquitous, the mass of notifications would quickly get too much to handle. This problematic tradeoff between providing awareness and overwhelming users was already reflected in related work [14]. When notices are not feasible, but the user is in control of the device, they can simply turn it off or unplug it as soon as they feel uncomfortable. Yet, whenever users do not have control options, researchers and developers can play an important role in developing assistive tools that give control back, such as done through the personalized privacy assistant [14], or smart home dashboards [67]. When the user can not exert control over the device, but the data collection is constant and ubiquitous, the only option left is to hand the user something that enables them to destroy or anonymize the data. Since such tools have to be explicitly leveraged by the user and are often characterized by unusual methods and appearance, we coined them guerilla tools. An existing example of such a tool is a wearable, developed by Chen et al. [12], that sends signals to disable microphones in the users' vicinity.

Whenever the data collection is not constant and ubiquitous, and the user has no control options, our decision tree states privacy notice as the most suitable approach (see Figure 4). Yet, our experts had conflicting opinions about that. While some strongly believe that providing awareness through informing the user about what is happening is the best option – especially when the user has no choice – other experts strongly opposed this idea as they believed it was wrong to distress users when they could not do anything. This leads to a fundamental question that HCI research alone cannot answer. *Should users be informed even when they do not have any options to limit or control a privacy violation? Or is it better not to inform users to prevent them from feeling powerless?* A possible way to approach this is by personalizing such notifications, for example, through a personalized privacy assistant [7]; this way, the notices would get adjusted to the individual's information need.

An important thing to consider is that the decision tree yields an idealistic point of view. It leads to the best privacy-preserving

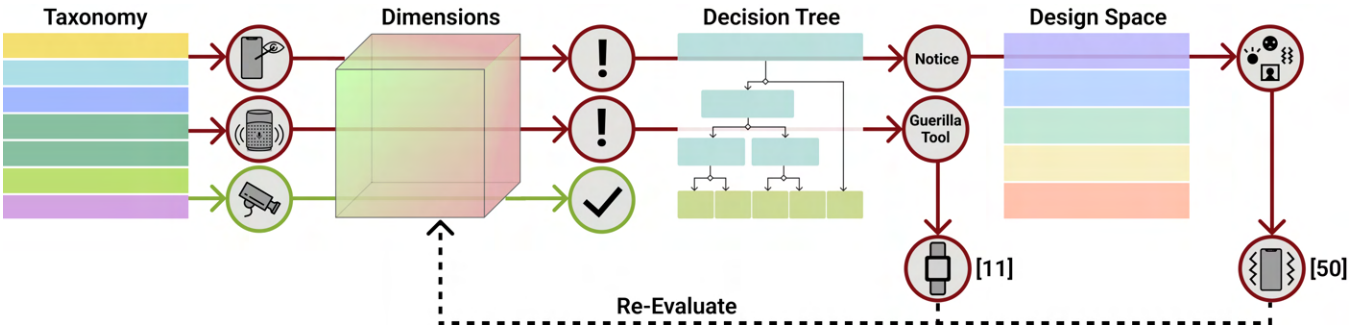

Figure 6: Three concrete examples (– shoulder surfing resulting in a notice as described in [12], – microphones resulting in a guerilla tool in the form of a bracelet as described in [53], - CCTV Camera resulting in no mitigation measure) of how our conceptual framework can be used to 1) create scenarios, 2) evaluate if they constitute a violation, 3) decide on the best mitigation measure, and 4), if applicable, design a notice.

mechanism possible, assuming the entity violating the privacy is interested in finding the best solution. However, we know that while privacy by design should be preferred whenever possible, many companies have a major interest in collecting their users' data un-anonymized. This trade-off was already brought up by Spiekermann and Cranor [59] that discuss that even though *privacy-by-architecture* (i.e., privacy by design) guarantees higher levels of privacy, many companies still rely on *privacy-by-policy* as it allows them to stick with their business model of collecting user data un-anonymized. Hence, businesses would be hesitant to obfuscate their data even though it might be technically possible. *Whenever this is the case, we need lawmakers to create regulations enforcing privacy protection in the best way possible – such as it is already done for the web.* Yet, companies often see users' data as compensation for providing a service for free. Another possibility would be offering an alternate compensation form, such as a monthly subscription for keeping users' data anonymized. This is, to some extent, already applied in practice. For instance, YouTube offers a premium subscription in exchange for not creating personalized profiles for targeted advertising.

6.3 A Conceptual Framework: Combining the Four Conceptual Tools

We provide four conceptual tools on technology-facilitated privacy violations in the physical world, which combined form a conceptual framework. First, our scenario taxonomy for privacy violations (see Figure 1) in the physical world can be used to identify the situation needing privacy-preserving mechanisms. As we know that different people have different perceptions of what constitutes a violation, our dimensions of privacy violations (see Figure 2) can clarify whether a situation is indeed violative. After identifying the situations needing a privacy-preserving mechanism, we can use the decision tree to decide on the best possible solution. While privacy by design should be preferred whenever possible, there are still situations where a notice (and choice) are the best available options. When we decide to use either a notice or notice and choice, our design space for privacy notices (see Figure 5) in the physical world can be used to explore when to deliver a notice, which media and modality should be used, and to decide on the content. Finally, after creating the notice, our dimensions can serve as an analysis tool to verify that a violation shifted successfully along the axes (see Figure 2) – at best, to the state of "voluntary and aware."

Figure 6 shows how our four conceptual tools play together based on three concrete examples. The first scenario ( ) depicts shoulder surfing and results in a privacy notice as described in the paper by Saad et al. [53]. First, we identify shoulder surfing as a relevant situation using the scenario taxonomy. We then use the dimensions, and, since the smartphone user does not voluntarily expose themselves to the shoulder surfing attack, conclude that it is indeed a violative situation. After that, we use the decision tree. As data obfuscation is impossible (as the user still wants to use their phone), the data collection is not constant and/or ubiquitous, and the user can not opt-out, we design a notice. We then use the design space to create four notification mechanisms: haptic feedback and three visual feedbacks. We then deploy the mechanisms and elicit their effectiveness by checking if they moved the violation along the dimensions to the state of voluntary and aware. The second scenario ( ) depicts the development of a guerilla tool in the form of a bracelet as a mitigation strategy for microphones, as described in the paper by Chen et al. [12]. We again use the taxonomy and identify a scenario where users are exposed to microphones. Using the dimensions, we conclude that a user is often forced and unaware of microphones being present, and thus, decide that it is a violative situation. We then use the decision tree, and as the data is not obfuscated, the collection is constant and ubiquitous (i.e., one is constantly surrounded by microphones, for example, through smartphones), and the person can not control these devices, decide to develop a guerilla tool. Finally, the last scenario ( ) depicts a situation where no mitigation measure was deemed necessary for a CCTV camera. We first create a scenario where a homeowner sets up a CCTV camera to surveil their garden. We then used the decision tree and concluded, as the homeowner set the camera up themselves, that they did it voluntarily and were aware of the data collection and possible consequences. As such, we decide that this situation is not a violation when judged from the homeowner's perspective and that no mitigation strategy is necessary.

6.4 Limitations

We used an online survey to collect experiences of privacy violations in the physical world which we clarified through in-depth interviews and design sessions with ten experts. An inherent drawback of our method of probing participants for experienced privacy violations is that we only capture those violations of which users are aware or at least become aware during some point of the interaction. Yet, as our goal was to capture which privacy-violating situations users experience during daily life, and not as previous work had already done (e.g., [22, 39]), to investigate perceptions towards scenarios defined by researchers, our taxonomy does not aim at being comprehensive in this regard. Yet, it will be important to compare and eventually combine the two perspectives in the future. Such an extended and encompassing taxonomy might especially aim at integrating situations of which users are unaware, as such situations will never be captured using our method of asking users. An extension might also aim at adding *why* the data collection happens as an additional layer. Users who find data collection beneficial might be more likely to consent. Another layer worth adding might be "who benefits from the data," as this might differ from the entity collecting the data and influence whether the collection is considered appropriate. Such an encompassing taxonomy might be used as one central point for retrieving scenarios needing privacy-preserving mechanisms and as the starting point for developing adequate measures.

While we hope to have obtained a comprehensive picture of user-experienced privacy-violating scenarios by covering different countries and, thus, cultural backgrounds, we cannot guarantee that we have covered all the possible scenarios. Moreover, with the growing number of smart devices and the ongoing expansion of the IoT, the taxonomy is expected to grow along all dimensions. Consequently, our taxonomy should not be seen as static but as a solid starting point, providing opportunities for future work to extend the taxonomy where applicable. Thus, we recommend repeating our online survey every 2 – 5 years to account for new technologies that have the potential to cause new violations through the fast advancement of the IoT.

We also want to discuss the possible term conflation of *real* and *physical*. We used both terms in our survey as we thought the term *real* might be more graspable for laypersons. We aimed at making clear to what kinds of violations we refer to by adding the notice about aiming at privacy violations in the *physical* world caused by technology and explicitly stating that we were not referring to privacy concerns caused by web pages or apps. However, technically, online privacy violations still happen in the *real* world. As such, the term conflation might be a reason for having to exclude 75 of our participants' scenarios. Consequently, future work should be consistent with terminology to exclude term conflation as a source of error.

The privacy-preserving mechanisms constructed in this investigation were, to some extent, limited to the capabilities of current technologies. Yet, we expect new technologies to shape our futures significantly through, for example, the widespread adoption of AR tools. This would enable new opportunities, not only for presenting but for actively engaging with privacy information directly in the environment. Hence, we also call for repeating the expert interviews to account for the technological innovations we expect in the upcoming years.

7 CONCLUSION

Through a large-scale online survey complemented by in-depth interviews and co-design sessions with privacy experts, we contributed four conceptual tools on technology-facilitated privacy violations in the physical world: (1) a scenario taxonomy based on experienced privacy violations in the physical world to identify situations in need of privacy-preserving mechanisms; (2) the dimensions of privacy violations to decide which situations are indeed a violation – as what constitutes a violation is very subjective. Additionally, the dimensions can help evaluate the effectiveness of implemented mechanisms; (3) a decision tree to decide on the best privacy-preserving mechanisms possible considering the individual characteristics of a situation; and finally, (4) a design space to explore designs for privacy notices in the physical world. Combined, our four tools build a conceptual framework to understand and mitigate technology-facilitated privacy violations in the physical world.

REFERENCES

[1] Noura Abdi, Xiao Zhan, Kopo M. Ramokapane, and Jose Such. 2021. Privacy Norms for Smart Home Personal Assistants. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 558, 14 pages. https://doi.org/10.1145/3411764.3445122

[2] Noah Apthorpe, Dillon Reisman, and Nick Feamster. 2017. A Smart Home is No Castle: Privacy Vulnerabilities of Encrypted IoT Traffic. *Workshop on Data and Algorithmic Transparency* abs/1705.06805 (2017). https://doi.org/10.48550/arXiv.1705.06805

[3] Natã M Barbosa, Joon S Park, Yaxing Yao, and Yang Wang. 2019. "What if?" Predicting Individual Users' Smart Home Privacy Preferences and Their Changes. *Proceedings on Privacy Enhancing Technologies* 2019, 4 (2019), 211–231. https://doi.org/10.2478/popets-2019-0066

[4] Florian Bemmann, Maximiliane Windl, Jonas Erbe, Sven Mayer, and Heinrich Hussmann. 2022. The Influence of Transparency and Control on the Willingness of Data Sharing in Adaptive Mobile Apps. *Proc. ACM Hum.-Comput. Interact.* 6, MobileHCI (2022), 26.

[5] Ann Blandford, Dominic Furniss, and Stephann Makri. 2016. *Qualitative HCI Research: Going Behind the Scenes.* Springer Cham, Cham, Switzerland. 51–60 pages. https://doi.org/10.2200/S00706ED1V01Y201602HCI034

[6] Frank Breitinger, Ryan Tully-Doyle, and Courtney Hassenfeldt. 2020. A Survey on Smartphone User's Security Choices, Awareness and Education. *Comput. Secur.* 88, C (jan 2020), 14 pages. https://doi.org/10.1016/j.cose.2019.101647

[7] Carnegie Mellon University. 2019. The Personalized Privacy Assistant Project. https://privacyassistant.org/

[8] Fred H Cate. 2010. The limits of notice and choice. *IEEE Security & Privacy* 8, 2 (2010), 59–62. https://doi.org/10.1109/MSP.2010.84

[9] Ann Cavoukian. 2009. Privacy by design: The 7 foundational principles. *Information and privacy commissioner of Ontario, Canada* 5 (2009), 2009.

[10] Supriyo Chakraborty and Omer Tripp. 2016. Eavesdropping and Obfuscation Techniques for Smartphones. In *Proceedings of the International Conference on Mobile Software Engineering and Systems* (Austin, Texas) *(MOBILESoft '16)*. Association for Computing Machinery, New York, NY, USA, 291–292. https://doi.org/10.1145/2897073.2897715

[11] Ankur Chattopadhyay and T.E. Boult. 2007. PrivacyCam: a Privacy Preserving Camera Using uCLinux on the Blackfin DSP. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, New York, NY, USA, 1–8. https://doi.org/10.1109/CVPR.2007.383413

[12] Yuxin Chen, Huiying Li, Shan-Yuan Teng, Steven Nagels, Zhijing Li, Pedro Lopes, Ben Y. Zhao, and Haitao Zheng. 2020. Wearable Microphone Jamming. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–12. https://doi.org/10.1145/3313831.3376304

[13] Richard Chow, Serge Egelman, Raghudeep Kannavara, Hosub Lee, Suyash Misra, and Edward Wang. 2015. HCI in Business: A Collaboration with Academia in IoT Privacy. In *HCI in Business*. Springer International Publishing, Cham, 679–687.

[14] 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.* Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3313831.3376389

[15] Lorrie Faith Cranor. 2022. Mobile-App Privacy Nutrition Labels Missing Key Ingredients for Success. *Commun. ACM* 65, 11 (oct 2022), 26–28. https://doi.org/10.1145/3563967

[16] Robert E. Crossler and France Bélanger. 2019. Why Would I Use Location-Protective Settings on My Smartphone? Motivating Protective Behaviors and the Existence of the Privacy Knowledge–Belief Gap. *Information Systems Research* 30, 3 (September 2019), 995–1006. https://doi.org/10.1287/isre.2019.0846

[17] Malin Eiband, Mohamed Khamis, Emanuel von Zezschwitz, Heinrich Hussmann, and Florian Alt. 2017. Understanding Shoulder Surfing in the Wild: Stories from Users and Observers. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) *(CHI '17)*. Association for Computing Machinery, New York, NY, USA, 4254–4265. https://doi.org/10.1145/3025453.3025636

[18] 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, New York, NY, USA, 447–464. https://doi.org/10.1109/SP40000.2020.00043

[19] Yuanyuan Feng, Yaxing Yao, and Norman Sadeh. 2021. A Design Space for Privacy Choices: Towards Meaningful Privacy Control in the Internet of Things. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 64, 16 pages. https://doi.org/10.1145/3411764.3445148

[20] Alisa Frik, Juliann Kim, Joshua Rafael Sanchez, and Joanne Ma. 2022. Users' Expectations About and Use of Smartphone Privacy and Security Settings. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 407, 24 pages. https://doi.org/10.1145/3491102.3517504

[21] Nina Gerber, Paul Gerber, Hannah Drews, Elisa Kirchner, Noah Schlegel, Tim Schmidt, and Lena Scholz. 2018. FoxIT: Enhancing Mobile Users' Privacy Behavior by Increasing Knowledge and Awareness. In *Proceedings of the 7th Workshop on Socio-Technical Aspects in Security and Trust* (Orlando, Florida, USA) *(STAST '17)*. Association for Computing Machinery, New York, NY, USA, 53–63. https://doi.org/10.1145/3167996.3167999

[22] Nina Gerber, Benjamin Reinheimer, and Melanie Volkamer. 2019. Investigating People's Privacy Risk Perception. *Proc. Priv. Enhancing Technol.* 2019, 3 (2019), 267–288. https://doi.org/10.2478/popets-2019-0047

[23] Nina Gerber, Verena Zimmermann, and Melanie Volkamer. 2019. Why Johnny Fails to Protect his Privacy. In *2019 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. IEEE, New York, NY, USA, 109–118. https://doi.org/10.1109/EuroSPW.2019.00019

[24] Colin M. Gray, Cristiana Santos, Nataliia Bielova, Michael Toth, and Damian Clifford. 2021. Dark Patterns and the Legal Requirements of Consent Banners: An Interaction Criticism Perspective. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 172, 18 pages. https://doi.org/10.1145/3411764.3445779

[25] Hana Habib, Megan Li, Ellie Young, and Lorrie Cranor. 2022. "Okay, Whatever": An Evaluation of Cookie Consent Interfaces. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 621, 27 pages. https://doi.org/10.1145/3491102.3501985

[26] Rakibul Hasan, Bennett I. Bertenthal, Kurt Hugenberg, and Apu Kapadia. 2021. Your Photo is so Funny That I Don't Mind Violating Your Privacy by Sharing It: Effects of Individual Humor Styles on Online Photo-Sharing Behaviors. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 556, 14 pages. https://doi.org/10.1145/3411764.3445258

[27] Rakibul Hasan, Yifang Li, Eman Hassan, Kelly Caine, David J. Crandall, Roberto Hoyle, and Apu Kapadia. 2019. Can Privacy Be Satisfying? On Improving Viewer Satisfaction for Privacy-Enhanced Photos Using Aesthetic Transforms. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3290605.3300597

[28] Hongxin Hu, Gail-Joon Ahn, and Jan Jorgensen. 2011. Detecting and Resolving Privacy Conflicts for Collaborative Data Sharing in Online Social Networks. In *Proceedings of the 27th Annual Computer Security Applications Conference* (Orlando, Florida, USA) *(ACSAC '11)*. Association for Computing Machinery, New York, NY, USA, 103–112. https://doi.org/10.1145/2076732.2076747

[29] Corey Brian Jackson and Yang Wang. 2018. Addressing The Privacy Paradox through Personalized Privacy Notifications. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2, 2, Article 68 (jul 2018), 25 pages. https://doi.org/10.1145/3214271

[30] Dilara Kekülüoğlu, Nadin Kökciyan, and Pınar Yolum. 2016. Strategies for Privacy Negotiation in Online Social Networks. In *Proceedings of the 1st International Workshop on AI for Privacy and Security* (The Hague, Netherlands) *(PrAISe '16)*. Association for Computing Machinery, New York, NY, USA, Article 2, 8 pages. https://doi.org/10.1145/2970030.2970035

[31] Dilara Kekülüoğlu, Nadin Kökciyan, and Pinar Yolum. 2018. Preserving Privacy as Social Responsibility in Online Social Networks. *ACM Trans. Internet Technol.* 18, 4, Article 42 (apr 2018), 22 pages. https://doi.org/10.1145/3158373

[32] Patrick Gage Kelley, Joanna Bresee, Lorrie Faith Cranor, and Robert W. Reeder. 2009. A "Nutrition Label" for Privacy. In *Proceedings of the 5th Symposium on Usable Privacy and Security* (Mountain View, California, USA) *(SOUPS '09)*. ACM, New York, NY, USA, Article 4, 12 pages. https://doi.org/10.1145/1572532.1572538

[33] Evan Lafontaine, Aafaq Sabir, and Anupam Das. 2021. Understanding People's Attitude and Concerns towards Adopting IoT Devices. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI EA '21)*. Association for Computing Machinery, New York, NY, USA, Article 307, 10 pages. https://doi.org/10.1145/3411763.3451633

[34] Josephine Lau, Benjamin Zimmerman, and Florian Schaub. 2018. Alexa, Are You Listening? Privacy Perceptions, Concerns and Privacy-Seeking Behaviors with Smart Speakers. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 102 (nov 2018), 31 pages. https://doi.org/10.1145/3274371

[35] Elena Maris, Timothy Libert, and Jennifer R Henrichsen. 2020. Tracking sex: The implications of widespread sexual data leakage and tracking on porn websites. *New Media & Society* 22, 11 (2020), 2018–2038. https://doi.org/10.1177/1461444820924632

[36] Karola Marky, Sarah Prange, Florian Krell, Max Mühlhäuser, and Florian Alt. 2020. "You Just Can't Know about Everything": Privacy Perceptions of Smart Home Visitors. In *19th International Conference on Mobile and Ubiquitous Multimedia* (Essen, Germany) *(MUM '20)*. Association for Computing Machinery, New York, NY, USA, 83–95. https://doi.org/10.1145/3428361.3428464

[37] Aleecia M McDonald and Lorrie Faith Cranor. 2008. The cost of reading privacy policies. *I/S: A Journal of Law and Policy for the Information Society* 4 (2008), 543–568. http://hdl.handle.net/1811/72839

[38] Andrés Molina-Markham, Prashant Shenoy, Kevin Fu, Emmanuel Cecchet, and David Irwin. 2010. Private Memoirs of a Smart Meter. In *Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building* (Zurich, Switzerland) *(BuildSys '10)*. Association for Computing Machinery, New York, NY, USA, 61–66. https://doi.org/10.1145/1878431.1878446

[39] 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

[40] Helen Nissenbaum. 2004. Privacy as contextual integrity. *Washington Law Review* 79 (2004), 119–158.

[41] Midas Nouwens, Ilaria Liccardi, Michael Veale, David Karger, and Lalana Kagal. 2020. *Dark Patterns after the GDPR: Scraping Consent Pop-Ups and Demonstrating Their Influence.* Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3313831.3376321

[42] Maggie Oates, Yama Ahmadullah, Abigail Marsh, Chelse Swoopes, Shikun Zhang, Rebecca Balebako, and Lorrie Faith Cranor. 2018. Turtles, locks, and bathrooms: Understanding mental models of privacy through illustration. *Proceedings on Privacy Enhancing Technologies* 2018, 4 (2018), 5–32.

[43] Jonathan A. Obar and Anne Oeldorf-Hirsch. 2020. The biggest lie on the Internet: ignoring the privacy policies and terms of service policies of social networking services. *Information, Communication & Society* 23, 1 (2020), 128–147. https://doi.org/10.1080/1369118X.2018.1486870

[44] Johannes Obermaier and Martin Hutle. 2016. Analyzing the Security and Privacy of Cloud-Based Video Surveillance Systems. In *Proceedings of the 2nd ACM International Workshop on IoT Privacy, Trust, and Security* (Xi'an, China) *(IoTPTS '16)*. Association for Computing Machinery, New York, NY, USA, 22–28. https://doi.org/10.1145/2899007.2899008

[45] Anne Oeldorf-Hirsch and Jonathan A. Obar. 2019. Overwhelming, Important, Irrelevant: Terms of Service and Privacy Policy Reading among Older Adults. In *Proceedings of the 10th International Conference on Social Media and Society* (Toronto, ON, Canada) *(SMSociety '19)*. Association for Computing Machinery, New York, NY, USA, 166–173. https://doi.org/10.1145/3328529.3328557

[46] Office of the California Attorney General. 2020. California Consumer Privacy Act (CCPA): Final Text of Proposed Regulations. https://gdpr.eu/cookies/

[47] Horizon 2020 Framework Programme of the European Union. 2021. Cookies, the GDPR, and the ePrivacy Directive. https://gdpr.eu/cookies/

[48] European Parliament. 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council. https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32016R0679

[49] Andrew S. Patrick and Steve Kenny. 2003. From Privacy Legislation to Interface Design: Implementing Information Privacy in Human-Computer Interactions. In *Privacy Enhancing Technologies*. Springer Berlin Heidelberg, Berlin, Heidelberg, 107–124. https://doi.org/10.1007/978-3-540-40956-4_8

[50] Robert W. Proctor, M. Athar Ali, and Kim-Phuong L. Vu. 2008. Examining Usability of Web Privacy Policies. *International Journal of Human-Computer Interaction* 24, 3 (2008), 307–328. https://doi.org/10.1080/10447310801937999
[51] Marianna Rapoport, Philippe Suter, Erik Wittern, Ondřej Lhóták, and Julian Dolby. 2017. Who You Gonna Call? Analyzing Web Requests in Android Applications. In *Proceedings of the 14th International Conference on Mining Software Repositories* (Buenos Aires, Argentina) *(MSR '17)*. IEEE Press, New York, NY, USA, 80–90. https://doi.org/10.1109/MSR.2017.11
[52] Joseph Reagle and Lorrie Faith Cranor. 1999. The Platform for Privacy Preferences. *Commun. ACM* 42, 2 (Feb. 1999), 48–55. https://doi.org/10.1145/293411.293455
[53] Alia Saad, Michael Chukwu, and Stefan Schneegass. 2018. Communicating Shoulder Surfing Attacks to Users. In *Proceedings of the 17th International Conference on Mobile and Ubiquitous Multimedia* (Cairo, Egypt) *(MUM 2018)*. Association for Computing Machinery, New York, NY, USA, 147–152. https://doi.org/10.1145/3282894.3282919
[54] Shruti Sannon, Brett Stoll, Dominic DiFranzo, Malte F. Jung, and Natalya N. Bazarova. 2020. "I Just Shared Your Responses": Extending Communication Privacy Management Theory to Interactions with Conversational Agents. *Proc. ACM Hum.-Comput. Interact.* 4, GROUP, Article 08 (jan 2020), 18 pages. https://doi.org/10.1145/3375188
[55] Florian Schaub, Rebecca Balebako, Adam L. Durity, and Lorrie Faith Cranor. 2015. A Design Space for Effective Privacy Notices. In *Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*. USENIX Association, Ottawa, 1–17. https://www.usenix.org/conference/soups2015/proceedings/presentation/schaub
[56] Helen Sharp, Jenny Preece, and Yvonne Rogers. 2019. *Interaction Design: Beyond Human-Computer Interaction*. Wiley. https://books.google.de/books?id=UDeQDwAAQBAJ
[57] Christopher Slobogin. 2002. Public privacy: camera surveillance of public places and the right to anonymity. *Miss. LJ* 72 (2002), 213.
[58] Daniel J. Solove. 2006. A Taxonomy of Privacy. *University of Pennsylvania Law Review* 154, 3 (2006), 477–564. https://doi.org/10.2307/40041279
[59] Sarah Spiekermann and Lorrie Faith Cranor. 2008. Engineering privacy. *IEEE Transactions on software engineering* 35, 1 (2008), 67–82. https://doi.org/10.1109/TSE.2008.88
[60] Eleftherios Spyromitros-Xioufis, Symeon Papadopoulos, Adrian Popescu, and Yiannis Kompatsiaris. 2016. Personalized Privacy-Aware Image Classification. In *Proceedings of the 2016 ACM on International Conference on Multimedia Retrieval* (New York, New York, USA) *(ICMR '16)*. Association for Computing Machinery, New York, NY, USA, 71–78. https://doi.org/10.1145/2911996.2912018
[61] Parth Kirankumar Thakkar, Shijing He, Shiyu Xu, Danny Yuxing Huang, and Yaxing Yao. 2022. "It Would Probably Turn into a Social Faux-Pas": Users' and Bystanders' Preferences of Privacy Awareness Mechanisms in Smart Homes. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 404, 13 pages. https://doi.org/10.1145/3491102.3502137
[62] Janice Y. Tsai, Serge Egelman, Lorrie Cranor, and Alessandro Acquisti. 2011. The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study. *Information Systems Research* 22, 2 (2011), 254–268. https://EconPapers.repec.org/RePEc:inm:orisre:v:22:y:2011:i:2:p:254-268
[63] Judith Odili Uchidiuno, Justin Manweiler, and Justin D. Weisz. 2018. Privacy and Fear in the Drone Era: Preserving Privacy Expectations Through Technology. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) *(CHI EA '18)*. Association for Computing Machinery, New York, NY, USA, 1–6. https://doi.org/10.1145/3170427.3188457
[64] Priyanka Verma and Sameer Patil. 2021. Exploring Privacy Aspects of Smartphone Notifications. In *Proceedings of the 23rd International Conference on Mobile Human-Computer Interaction* (Toulouse & Virtual, France) *(MobileHCI '21)*. Association for Computing Machinery, New York, NY, USA, Article 48, 13 pages. https://doi.org/10.1145/3447526.3472065
[65] Rolf H. Weber. 2010. Internet of Things – New security and privacy challenges. *Computer Law & Security Review* 26, 1 (2010), 23–30. https://doi.org/10.1016/j.clsr.2009.11.008
[66] Maximiliane Windl, Niels Henze, Albrecht Schmidt, and Sebastian S. Feger. 2022. Automating Contextual Privacy Policies: Design and Evaluation of a Production Tool for Digital Consumer Privacy Awareness. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 34, 18 pages. https://doi.org/10.1145/3491102.3517688
[67] Maximiliane Windl, Alexander Hiesinger, Robin Welsch, Albrecht Schmidt, and Sebastian S. Feger. 2022. SaferHome: Interactive Physical and Digital Smart Home Dashboards for Communicating Privacy Assessments to Owners and Bystanders. *Proc. ACM Hum.-Comput. Interact.* 6, ISS, Article 586 (nov 2022), 20 pages. https://doi.org/10.1145/3567739
[68] Maximiliane Windl and Sven Mayer. 2022. The Skewed Privacy Concerns of Bystanders in Smart Environments. *Proc. ACM Hum.-Comput. Interact.* 6, MobileHCI (2022), 21. https://doi.org/10.1145/3546719
[69] Xiaokui Xiao and Yufei Tao. 2006. Personalized Privacy Preservation. In *Proceedings of the 2006 ACM SIGMOD International Conference on Management of Data* (Chicago, IL, USA) *(SIGMOD '06)*. Association for Computing Machinery, New York, NY, USA, 229–240. https://doi.org/10.1145/1142473.1142500
[70] Yaxing Yao, Justin Reed Basdeo, Smirity Kaushik, and Yang Wang. 2019. Defending My Castle: A Co-Design Study of Privacy Mechanisms for Smart Homes. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–12. https://doi.org/10.1145/3290605.3300428
[71] Yaxing Yao, Justin Reed Basdeo, Oriana Rosata Mcdonough, and Yang Wang. 2019. Privacy Perceptions and Designs of Bystanders in Smart Homes. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 59 (nov 2019), 24 pages. https://doi.org/10.1145/3359161
[72] Yaxing Yao, Huichuan Xia, Yun Huang, and Yang Wang. 2017. Privacy Mechanisms for Drones: Perceptions of Drone Controllers and Bystanders. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) *(CHI '17)*. Association for Computing Machinery, New York, NY, USA, 6777–6788. https://doi.org/10.1145/3025453.3025907

A SURVEY

A.1 Demographics

(1) In which country do you currently reside? (drop-down list)
(2) Which gender do you most identify with?
  - Male
  - Female
  - Non-binary
  - Self-described
(3) How old are you? (number field)
(4) What is the highest degree you have received?
  - Less than high school degree
  - High school graduate
  - Some college but no degree
  - Bachelor's degree
  - Master's degree
  - Doctoral degree
  - Vocational education
(5) What is your current primary occupation? (free text)
(6) Would you describe yourself as a "privacy expert"? (yes/no)

A.2 Questionnaire

Privacy is important to all of us. Therefore, humans tend to communicate their personal information selectively. However, due to the prevalence of technology, our surroundings (e.g., people, devices, companies, governments) continuously try to collect and infer information about us. In such cases, we do not know what information is collected or by whom. This also means we can not give active consent. As a result, we might feel like our privacy has been violated. Note: We aim to cover privacy violations in the physical world caused by technology. This means we are **not referring to privacy concerns caused by using web pages or apps.**

(1) Have you ever felt like your privacy was violated in the real world by means of technology? (yes/no)
(2) What type of privacy concerns does this survey address? (sanity check)
  - Privacy violations in the real world
  - Privacy violations when using web pages/apps

  *[If answered **yes** to question (1)]*
(3) For the situations in which you have felt that your privacy was violated by technology, describe the scenarios as concretely as possible, including:
  1) where you were when the violation happened,

2) which technology caused the violation, and
3) what kind of private information was affected.
Provide at least 3 scenarios. If you encountered less than 3, envision new scenarios. Please keep in mind that we are **not referring to privacy concerns caused by using web pages or apps.** (free text)

(4) Did you experience the scenario or envision it?
   (a) I experienced it.
   (b) I envisioned it.
   *[If answered **no** to question (1)]*

(5) Envision real-world scenarios where your privacy can be violated by technology. Describe the scenarios as concretely as possible, including:
1) where you were when the violation happened,
2) which technology caused the violation, and
3) what kind of private information was affected.
Provide at least 3 scenarios. If you encountered less than 3, envision new scenarios. Please keep in mind that we are **not referring to privacy concerns caused by using web pages or apps.** (free text)

(6) This is the last question of the survey. If you have any additional feedback, please let us know here. (free text)