

Decide Yourself or Delegate – User Preferences Regarding the Autonomy of Personal Privacy Assistants in Private IoT-Equipped Environments

Karola Marky
Ruhr University Bochum
Bochum, Germany
Technical University of Darmstadt
Darmstadt, Germany
karola.marky@rub.de

Kira Bleck
Paul Gerber
Technical University of Darmstadt
Darmstadt, Germany

Florian Alt
University of the Bundeswehr Munich
Munich, Germany
florian.alt@unibw.de

Alina Stöver
Technical University of Darmstadt
Darmstadt, Germany
stoever@psychologie.tu-
darmstadt.de

Verena Zimmermann
ETH Zürich
Zürich, Switzerland
verena.zimmermann@gess.ethz.ch

Max Mühlhäuser
Technical University of Darmstadt
Darmstadt, Germany
max@informatik.tu-darmstadt.de

Sarah Prange
University of the Bundeswehr Munich
Munich, Germany
sarah.prange@unibw.de

Florian Müller
LMU Munich
Munich, Germany
florian.mueller@um.ifi.lmu.de

ABSTRACT

Personalized privacy assistants (PPAs) communicate privacy-related decisions of their users to Internet of Things (IoT) devices. There are different ways to implement PPAs by varying the degree of autonomy or decision model. This paper investigates user perceptions of PPA autonomy models and privacy profiles – archetypes of individual privacy needs – as a basis for PPA decisions in private environments (e.g., a friend’s home). We first explore how privacy profiles can be assigned to users and propose an assignment method. Next, we investigate user perceptions in 18 usage scenarios with varying contexts, data types and number of decisions in a study with 1126 participants. We found considerable differences between the profiles in settings with few decisions. If the number of decisions gets high ($> 1/h$), participants exclusively preferred fully autonomous PPAs. Finally, we discuss implications and recommendations for designing scalable PPAs that serve as privacy interfaces for future IoT devices.

CCS CONCEPTS

• **Security and privacy** → **Privacy protections**; • **Human-centered computing** → *Empirical studies in HCI*.

KEYWORDS

privacy, IoT, privacy profiles, personal privacy assistance

ACM Reference Format:

Karola Marky, Alina Stöver, Sarah Prange, Kira Bleck, Paul Gerber, Verena Zimmermann, Florian Müller, Florian Alt, and Max Mühlhäuser. 2024. Decide Yourself or Delegate – User Preferences Regarding the Autonomy of Personal Privacy Assistants in Private IoT-Equipped Environments. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*, May 11–16, 2024, Honolulu, HI, USA. ACM, New York, NY, USA, 20 pages. <https://doi.org/10.1145/3613904.3642591>

1 INTRODUCTION

With the increasing spread of IoT devices in private households [66], we no longer only encounter such devices in our own homes, but also in all kinds of semi-public contexts such as when visiting friends [45, 78] or in rental holiday homes [43, 65, 71]. While we can configure the private devices we control according to our personal privacy preferences, we have no control over the settings of the devices we encounter in the mentioned scenarios. Even worse, we have no access to the configuration of such devices and do not know what private information they record and process about us.

The level of desired privacy is highly individual: it varies among people and cultures [7]. Hence, individuals need the power to decide about data sharing in a world that increasingly “runs” on captured data, for example, by having automated processes at home. Yet, IoT devices and their data collection capabilities are not standardized which poses a challenge to privacy controls. For instance, one smart TV might not collect any data while another one has integrated cameras and microphones. How can individuals judge that without knowing the specific TV model? What about “hidden” IoT devices that are seamlessly integrated into a living space? Most devices are controlled remotely and it may be challenging to assess their current status [13]. To ultimately empower individuals, researchers came up with Personalized Privacy Assistants (PPAs) [16] – trusted agents that can notify users about IoT devices in their environment,

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
CHI '24, May 11–16, 2024, Honolulu, HI, USA
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0330-0/24/05.
<https://doi.org/10.1145/3613904.3642591>

provide them with decision support, or even make a decision on behalf of their users [13].

In the context of this paper, PPAs are considered assistive software that runs on a personal device of the user. All IoT devices capturing personal data, send data capturing requests to the users through the PPA. Furthermore, a PPA can help its users discover IoT devices in foreign places, such as rental holiday homes, hotel rooms, or the flat of a friend. To become trusted agents, PPAs need knowledge about their users: they have to know to which degree users wish to be involved in the decision and their data-sharing preferences (e.g., considering the user's context, such as their location).

Related work has investigated several aspects of PPAs, including how they can technically predict or model privacy decisions [3, 20, 31, 81], how people might be segmented into groups based on privacy preferences [17, 34, 41], and how individuals want to communicate with a PPA [13]. In this context, a large-scale analysis of mobile privacy behavior has shown that privacy profiles might be a viable solution to simplify how users grant access to their data [41].

We integrate the body of knowledge from related work to explore autonomous PPAs. First, we investigate privacy profiles as a basis for autonomous PPAs with the potential to significantly reduce user effort, following a multi-step approach. Privacy profiles summarize the privacy needs of users through archetypes. The first research question is:

RQ1: *How can privacy profiles be used to create autonomous PPAs?*

We use the five well-established privacy profiles from Dupree et al. [17] as a basis, to propose a possible extension and a short questionnaire for assigning users to a profile. For this, we conducted a series of three user studies, totaling $N = 417$ participants, to develop and validate the profile assignment questionnaire and the final set of privacy profiles. Then, we investigated different PPA autonomy models in 18 scenarios in the main study ($N = 1126$), investigating:

RQ2: *Which PPA decision support types are preferred?*

RQ3: *How do users' wishes for PPA support differ based on the context of use?*

To answer these research questions, we conducted an online study, introducing three PPA autonomy types, 1) asking users for each decision (notification PPA), 2) recommending a decision to the user (recommendation PPA), and 3) making a decision for the user (fully autonomous PPA), based on Colnago et al.'s work [13]. Participants rated the autonomy type preference across 18 usage scenarios that differed with regards to a) the location the IoT device was located, b) the type of data captured and c) the number of decisions that refer to the privacy decisions users have to make.

We show that individuals' privacy preferences, i.e., different privacy profiles, play a more significant role in foreign environments where the number of privacy decisions is low. Here, participants with profiles that describe high motivations to protect their privacy preferred notification PPAs. All other privacy profiles either preferred recommendations or fully autonomous PPAs. This was also true for IoT devices located in foreign households because the home and IoT devices in it are typically well-known. However,

these preferences shift when the number of decisions increases: here, all privacy profiles converge to fully autonomous PPAs because the burden on users is very high. While we investigated PPAs for private IoT-equipped environments, our assignment method for a privacy profile may also be valuable for other domains where clustering users based on privacy needs is useful, such as sharing data on social media.

Contribution Statement. This paper presents a comprehensive investigation of PPA autonomy types for IoT devices through large-scale investigations. Our contribution is threefold:

- (1) We propose a **set of privacy profiles** based on the work by Dupree *et al.* [17] alongside a possible assignment method as a basis for PPA decisions in private IoT-equipped environments.
- (2) We investigate aspects of **PPA autonomy** in 18 scenarios focused on the own home and the home of a friend. Based on the results, we discuss challenges and future research for designing and implementing scalable PPAs.
- (3) This work leverages the results of existing PPA studies in other domains (cf. [19, 41]) and extends qualitative work on decision and autonomy models of PPAs to different scenarios [13] while also considering scalability challenges when communicating with IoT devices [47].

2 BACKGROUND AND RELATED WORK

In this section, we first summarize studies that investigated the *privacy perceptions of IoT devices*. Then, we describe research that investigated *personalized privacy assistants*. Finally, we detail existing *privacy profiles*.

2.1 Privacy Perceptions of IoT Devices

Many researchers investigated privacy attitudes in the scope of IoT devices and identified privacy concerns [1, 2, 4, 10–12, 67, 69, 71, 75, 76, 80]. In summary, related work has shown that users and bystanders in IoT environments require assistance in the tasks of making, formulating, communicating, and executing their privacy decisions. In this work, we propose a concept that serves as a basis for a PPA that assists users and bystanders in different tasks related to privacy decisions based on their personal preferences.

Investigations of specific privacy attitudes have produced mixed results. On one hand, participants in studies feared that the collected data might be misused to plan burglaries or private information might be assessed by unwanted third parties [79, 83, 84]. This aspect has been confirmed by studies of IoT device owners and further participants who were not concerned about the collected data but did not trust IoT device providers and wished for transparency [19, 29, 51, 57, 67]. Yao *et al.* let participants design privacy-respecting IoT interfaces, resulting in designs that increased the transparency of data collection and allowed controlling the data collection [77]. Users are generally willing to share even privacy-sensitive data with service providers if the data cannot be linked to them [35]. On the other hand, investigations are showing that users are not concerned about potential threats [37, 79]. These mixed results might be related to differences in participants' knowledge, their culture, differences regarding the IoT devices they used, and different privacy attitudes.

Several studies have focused on special user groups, such as bystanders, that neither buy nor control the devices. This could be minors or other inhabitants in smart homes [50, 70] or smart home visitors [44, 47, 74, 78]. Within this scope, minors might feel discomfort, for example, as parents can monitor them [49, 56, 70]. A co-design study of three scenarios in which the privacy of bystanders might be affected revealed that bystanders wish to exert control over data collected about them [78]. Yet, there are hurdles on the social level [44, 74]. Further investigations of bystander privacy in IoT environments concluded that scalability constitutes a major challenge for supporting bystander privacy [46, 47].

Rental holiday homes, such as AirBnBs, were investigated as a special use case. Here, research highlights that users either need support in discovering devices [43, 65] or the information should be available on the booking website [71].

Another stream of research investigated how privacy in the presence of IoT devices could be improved. Being aware of potential consequences motivated participants to configure IoT devices matching their privacy needs [25, 33]. Privacy settings interfaces and the information provided by them have been investigated in this scope, showing that users wish for access to detailed information for making a decision [48]. Another possibility to enhance privacy is informing users about IoT devices through labels on the device packaging before they buy them [20].

2.2 Personalized Privacy Assistants (PPA)

In general, *privacy* means that users can control the circumstances and conditions under which personal information about them is collected and processed by a third party [15, 72]. As a result, each user individually decides about private data. To do so, users first have to evaluate information to *make* the decision. Then, they *formulate* it to express their intention. Finally, they *communicate* the decision to another person or entity that can *realize it*. Users can be overwhelmed by privacy decisions for several reasons: they might be overwhelmed by the information that has to be considered for making a decision [13] or by the number of daily decisions [47].

Personalized Privacy Assistants (PPAs) are systems that assist users in protecting their privacy. Depending on the specific implementation of the PPA, it can support the users in the different tasks of making, formulating, communicating, and realizing privacy decisions. PPAs have been investigated in different domains, such as websites [14], online social media [22, 28], or mobile apps [40, 68]. In the remainder of this section, we focus on PPAs for IoT devices.

Colnago *et al.* describe three possibilities for realizing PPAs: 1) notification PPAs inform users about data collection, 2) recommendation PPAs provide decision recommendations, and 3) decision PPAs act on behalf of the user [13]. An investigation of these three PPA implementations shows that users, in general, prefer PPAs with control possibilities, which are considered more important than cognitive overload [13]. Specific scenarios were not investigated. Existing work on PPAs for IoT devices strongly focused on public environments. For instance, Pappachan *et al.* proposed a framework for PPAs in smart buildings [53]. Langheinrich describes a PPA framework, informing users about the data collection and providing access to privacy settings [36]. Das *et al.* developed an app that informs users about nearby cameras [16]. Raber *et al.*

addressed the data collection problem by smart retail stores, like Amazon Go, by building a privacy manager that can help users with their privacy settings [55]. Naeini *et al.* conducted a large-scale online survey, showing that privacy attitudes are connected to environments and data types [19]. People perceived data collection in public environments as less critical than in private ones. Such private environments can be smart homes. He proposes a privacy settings interface that allows users to configure multiple smart assistants [26]. Seymore extended this idea and developed a PPA that not only informs users about the data collection of their IoT devices but also provides the option to set up a firewall to prevent data leakage and provides lessons about network privacy [63].

Our work builds upon these existing works. We investigate the PPA implementations from Colnago *et al.* [13] and how users would like to adapt them in specific usage contexts. These contexts are informed by the results from Naeini *et al.* [19] including the scalability challenge for PPAs by Marky *et al.* [47]. Naeini *et al.* used an ML approach for decision-making. In contrast, we investigate using a privacy profile as a basis. The different realizations of the PPA allow for different user engagement as they support tasks of the decision-making and communication process.

2.3 Privacy Profiles

We investigate privacy profiles as a basis for PPAs. A privacy profile describes the privacy preferences of a group of users. Those preferences are clustered based on privacy attitudes. Liu *et al.* [41] showed that a small number of clusters, that can be used as privacy profiles, results from analyzing user behavior on (not) granting app permissions. In this paper, we consider the data-capturing requests from IoT devices as similar to permission requests from apps. While Liu *et al.* considered a very narrow context, several broader approaches for clustering users into profiles are proposed in the literature. These vary from a relation to a specific technology, such as the IUIPC covering Internet usage [42], to technology-independent approaches, such as the concern for information privacy (CFIP) [64], representing privacy concerns as numerical values. On the other hand, approaches for profile clustering are connected to different constructs. For instance, the information-seeking preferences by Morton *et al.* [52] are related to information-seeking.

An often-cited approach, unrelated to a specific technology, stems from Westin, who assigns users based on their privacy concerns into three clusters [34]. Other approaches cluster users with a focus on their privacy needs in certain application areas, like mobile app permissions [39]. Dupree *et al.* [17] classify individuals according to their privacy knowledge and motivation to protect their privacy. This paper is based on the privacy profile from Dupree *et al.* since those are related to technical systems without being too specific. Furthermore, prior investigations of Dupree's privacy profiles showed that developers found it helpful to use the profiles as a basis to design for certain user types [60]. The five privacy profiles from Dupree *et al.* [17] are as follows¹:

¹Note, that the titles of the profiles were proposed by Dupree *et al.* [17], yet use non-ideal language as terms like "lazy" might be perceived as overly judgmental. We use the titles from Dupree *et al.* in the related work, method, and analysis parts of this paper and propose more neutral and inclusive language in the discussion.




Figure 1: In three pre-studies ($N = 417$), we investigated privacy profiles and developed an assignment method. In the main study ($N = 1126$), we investigated choices of implementations for personal privacy assistance in different scenarios.

- (1) *Marginally Concerned* (low knowledge and motivation): These users have little knowledge about privacy protection mechanisms and do not fear cyber attacks. Hence, they are not motivated to extend their knowledge.
- (2) *Struggling Amateur* (medium knowledge and motivation): These users have knowledge about privacy protection and limit the information they share with others. Even if their motivation is limited, they prefer to protect their privacy if they receive information to do so.
- (3) *Lazy Expert* (high knowledge, low motivation): Lazy Experts have detailed knowledge about privacy protection. However, they prefer convenience over privacy.
- (4) *Fundamentalist* (high knowledge, high motivation): Fundamentalists have detailed knowledge, similar to lazy experts. Their motivation drives them to help others protect their privacy.
- (5) *Technician* (medium knowledge, high motivation): Technicians are highly motivated and are aware of possible consequences. However, their knowledge is lower compared to lazy experts and fundamentalists.

In contrast to existing work, we investigate how the privacy profiles from Dupree *et al.* can be leveraged as a basis for a PPA. Hence, we develop a method for profile assignment and investigate preferences in different usage scenarios that are connected to the assigned privacy profile.

3 PERSONALIZED PRIVACY ASSISTANCE (PPA) BASED ON PRIVACY PROFILES

In this section, we detail PPA based on privacy profiles. We propose a PPA concept and investigate methods for *assigning users to a privacy profile*. There are several ways of implementing PPAs. Those differ regarding the specific tasks PPAs perform for their users, decision-making strategies, and PPA autonomy. Previous studies have shown that users' privacy attitudes differ from each other [80]. Thus, users wish for different types of PPAs. As shown by Colnago *et al.*, users generally prefer to be in control when using a PPA [13]. Yet, the degree of preferred control might differ among users. We propose the following concept for a PPA where the PPA performs tasks for the user based on their privacy profile. This profile describes the user's privacy attitude and the tasks the user wishes support for.

Considering the usage of this PPA, users first have to be *assigned* to a privacy profile during the PPA setup. This assignment is based on a short questionnaire, as suggested by Rudolph *et al.* [60]. Next, the user receives a summary of the privacy support by the specific

privacy profile containing the tasks that the assistant does for the user. During usage, the PPA runs on the user's mobile device as a trusted agent. The user can access logs of the PPA tasks anytime and update PPA settings. We also consider that the user can overwrite properties of a privacy profile. In the next section, we describe the iterative development of a questionnaire for assigning users to one of the privacy profiles by Dupree *et al.* [17]. This assignment serves as a basis for our PPA but can also be used in other domains.

3.1 Profile Assignment

To investigate privacy assistance based on privacy profiles, we started with a series of three pre-studies with a total of 417 participants. During these three pre-studies, we developed a method for assigning users to a privacy profile. Above all, this method should be time-efficient so that it can be realistically implemented in a PPA. Figure 1 provides an overview of our study approach.

3.1.1 Pre-Study I: Investigation of Dupree Profiles. As a starting point for profile assignment, we investigated existing privacy profiles from Dupree *et al.* [17] detailed above in an exploratory study.

Goal. The goal of our first step was a quantitative investigation of the privacy profiles from Dupree *et al.* [17]. These profiles were empirically motivated through qualitative studies and validated by experts. Consequently, we investigate how well the profiles fit user attitudes.

Captured Data. The profiles are based on the two dimensions 'knowledge about privacy protection' and 'motivation to protect one's privacy'. We designed two simple items to capture these dimensions. We opted for two items on an ordinal scale to enable a time-efficient assessment. We asked participants to self-report their levels of knowledge and motivation on three levels (low, medium, and high). The levels were chosen based on the profile descriptions from Dupree *et al.* [17]. The items were: (1) How would you rate your privacy knowledge? and (2) How would you rate your motivation to protect your information privacy?

Participants. We recruited 67 participants via mailing lists. Of those, 31 identified as female, 35 as male, and one preferred not to specify. The sample's mean age was $M=29.27$ ($SD=11.12$). A total of 46 participants held a university or college degree or a PhD, 15 completed A-levels, three had another school-leaving qualification, and four had other qualifications, such as a completed apprenticeship.

Results & Takeaways. For the evaluation of the study results, we first constructed a 3×3 matrix, describing the existing privacy profiles

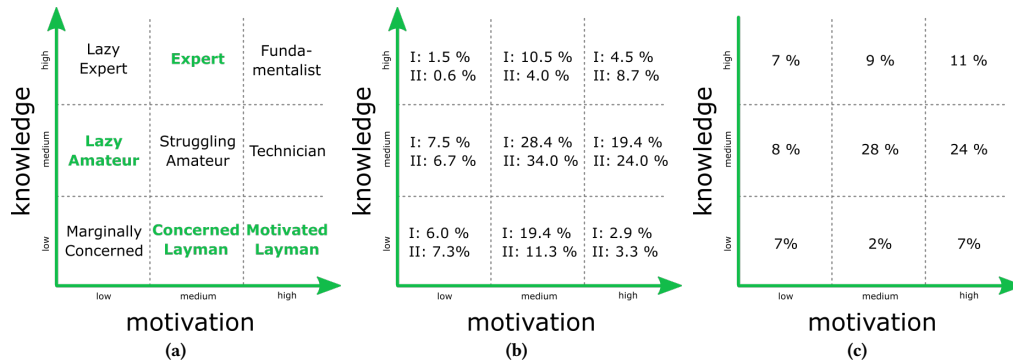


Figure 2: (a) Depiction of privacy profiles, our extension is marked in green, (b) distribution of profiles in pre-study I and II, and (c) distribution of profiles in pre-study III.

by Dupree (see also Fig. 2a). There are nine possible combinations of knowledge and motivation since each scale has three possible values. However, the existing profiles from Dupree *et al.*, cover five spots of this matrix. From our 67 participants, 40 could be assigned to an existing privacy profile from Dupree (see Figure 2b). Knowledge and motivation of the remaining 27 participants were located in the four spots that were not covered by a privacy profile. Two reasons can explain these results. Either the profiles by Dupree *et al.* oversimplify privacy attitudes and, hence, do not cover the complete spectrum; or the self-reporting of knowledge and motivation is unreliable. We conducted a follow-up study to investigate these two aspects in more detail.

3.1.2 Pre-Study II: Questionnaire and Profile Development. In our second pre-study, we investigated alternatives to self-reporting and constructs to describe the privacy profiles.

Goal. The study goal was investigating alternatives to self-reporting of knowledge and motivation by identifying constructs suitable to describe the privacy profiles in more detail. Pre-study I found clusters of privacy attitudes that are not covered by the existing profiles. Hence, we investigated to which extent the clusters exist.

Identifying Scales & Constructs. To build a theoretical foundation for targeted investigation of privacy clusters among users, we reviewed the literature to identify suitable constructs associated with privacy behaviors and attitudes. In cooperation with privacy research experts, the work on privacy indexes by Kumaraguru and Cranor [34] was chosen as a starting point for a snowball search. To not miss further import work that did not cite Kumaraguru and Cranor, we used Google Scholar as search space with the search terms *Privacy Personas*, *Privacy Paradox*, and *Privacy*, limited to Human-Computer Interaction conferences. Each identified paper was read to extract constructs associated with privacy behaviors and attitudes. In this step, we also considered constructs that do not measure privacy directly but have items connected to it. The literature search yielded eleven constructs, provided in Appendix A. Some of these theoretical constructs had already been identified as useful for forming clusters through prior research, while others were subject to exploratory investigation. To not randomly limit

the validity of our results, we opted to use all eleven constructs in addition to the self-formulated “Privacy Motivation and Knowledge” construct consisting of the two items from pre-study I in this step. Thus, twelve constructs were evaluated overall.

Participants. We conducted another online study in which we presented all items to 150 participants recruited via mailing lists. Of these, 95 identified as female, 54 as male, and one preferred not to specify. The sample’s mean age was $M=33.12$ ($SD=14.2$). A total of 98 participants held a university degree, 39 completed A-levels, seven had other school-leaving qualifications, and six had other qualifications, such as a completed apprenticeship.

Results & Takeaways. We first explored the 3x3 matrix resulting from the responses to the privacy knowledge and motivation questions. The distribution of the participants among the different levels of knowledge and motivation was similar to the first study (see Fig. 2b). In addition to this matrix, we assessed the responses to the questionnaires covering an additional eleven constructs as listed in Appendix A. We created graphical plots of the means of the responses to the different questionnaires and items for a visual inspection. Next, two researchers exploratorily identified items for which the mean responses of the participants in different matrix cells formed patterns that could be linked to the profile clusters by Dupree *et al.* and the four empty spots. The aim was to select items for a meaningful distinction regarding participants’ privacy knowledge and motivation. Using this method, we selected nine specific items (see Appendix A.2) from the twelve constructs (see Appendix A) that can be used to complement the two knowledge and motivation questions to distinguish the nine clusters of privacy attitudes better. Furthermore, the item content was used to develop initial descriptions of the clusters not covered by the Dupree *et al.*’s privacy profiles. Fig. 2a shows the extension together with the Dupree profiles. The developed four descriptions are as follows:

- (1) *Concerned Layman* (low knowledge and medium motivation): Their privacy is important to them, and they are motivated to protect it. At the same time, they do not have knowledge about privacy protection and are not interested in the details of how to do so (e.g., they are not interested in reading articles about security threats).

- (2) *Motivated Layman* (low knowledge and high motivation): Users in this profile are highly motivated to protect their privacy but do not know how to. Privacy is important to them, but they are not always interested in educating themselves on this topic. Furthermore, they do not make quick decisions.
- (3) *Lazy Amateur* (medium knowledge and low motivation): Lazy Amateurs are not motivated to protect their privacy, even though it is important to them and they have some knowledge of how to do so. Usually, they make quick decisions but often procrastinate when it comes to making important decisions.
- (4) *Expert* (high knowledge and medium motivation): Privacy is important to them, and they have high knowledge but have medium motivation to protect it. It is also essential to them that they are aware of and knowledgeable about how their personal information will be used.

After completing pre-study II, we had a complete description of the privacy attitude clusters and a set of items for the profile assignment. Furthermore, through this study, we collected mathematical data in terms of nine profile vectors that describe the nine privacy profiles in the Euclidean space. Based on that, we used the answers to the nine questionnaire items to assign a profile to a participant mathematically. Two essential aspects were missing at this point: (1) the assigned profiles might not match reality, and (2) the Euclidean space is uniform and does not consider potential distortions. To (1) check whether the assigned profiles indeed match the attitudes of users and (2) collect data to account for the distortions mathematically, we conducted pre-study III.

3.1.3 Pre-Study III: Questionnaire Evaluation. With our third and final pre-study we evaluated and validated the developed assignment questionnaire.

Goal. To validate the questionnaire, we (a) assigned participants to a profile and let them evaluate the profile fit and (b) collected data to improve our assignment method.

Participants. We recruited a sample of 200 participants from the Prolific online platform. Of these, 66 identified as female, 131 as male, two as other, and one preferred not to specify. The sample's mean age was $M=27.32$ ($SD=8.92$). A total of 114 participants held a university degree or Ph.D., 69 completed A-levels, thirteen were school students, and four reported other qualifications.

Study Procedure. First, we asked participants to fill in our profile assignment questionnaire (see Appendix A.2). Then, we assigned a privacy profile by comparing the answer vector from the participants to a vector of the privacy profile. We measured the distance between the two vectors using cluster analysis with the Euclidean distance and presented the closest profile to the participant. Cluster analysis is used to group investigated objects in such a way that the differences between the objects of a group (in our case the privacy profiles) are as small as possible and the differences between the clusters are as large as possible [9].

Once the profile assignment was completed, participants were shown the description of their assigned profile (no profile title). We then asked them to rate how well the profile matches their privacy

preferences on a 5-point Likert scale (1="not at all"; 5="excellent") and to explain their choice in a text field. If the rating was 3 or lower, we presented the descriptions of all nine privacy profiles consecutively and asked them to rate the fit of each profile and to choose the profile they thought described them best.

Results & Takeaways. Overall, participants rated the profile fit with 3.83 ($SD = 0.88$). In total, 60 participants (30%) rated the fit as three or lower and thus assigned themselves another profile, which resulted in a final profile fit of 4.19 ($SD = 0.60$). Fig. 2c shows the distribution of the profiles within our sample. While the distribution overall appears to be similar to the other pre-studies, we observed differences regarding the profile "concerned layman". The differences might be explained through the refined assignment method in which participants themselves evaluated the fit with the assigned profile. Consequently, participants have been reassigned to another profile that better fits their attitudes.

As stated above, profile assignment by the Euclidean distance neglects distortions and correlations in the vector space. For the assignment in pre-study III, we used the Euclidean distance because we did not have any information about correlations in the data set. Using the data collected in pre-study III, we refined our assignment method by calculating a covariance matrix. Hence, in the main study (see Section 4), we used the Mahalanobis distance. For each privacy profile, we calculated the mean vector of the nine profiles and the corresponding covariance matrix².

4 MAIN STUDY METHODOLOGY

To investigate whether the extended privacy profiles can serve as a basis for a PPA, we conducted an online study, recruiting $N = 1126$ participants from Prolific. Based on the distribution of the privacy profiles in our pre-studies, we opted for an online study to collect a large sample and enough responses for each profile. Before the actual study, we ran a pilot with five experts and another one with ten participants to improve the clarity of questions and instructions.

4.1 Survey Procedure

The procedure of our online survey was as follows.

4.1.1 Welcome and Consent. First, participants received the consent form and were asked to read and accept it.

4.1.2 Experience, Understanding of IoT, and Anchoring. In this part, we captured the participants' previous experiences with IoT devices. We asked them whether they had already heard about the Internet of Things and how. Next, we normalized the participants' knowledge of IoT devices to make sure that all participants shared a common baseline. We presented them with a description of the IoT and asked which IoT devices they owned. The recruitment platform also offers a filtering option to identify Prolific users who are familiar with IoT devices. However, previous studies (cf. [44]) showed that this function is not reliable which is why, we let participants report their experiences.

²One privacy profile, concerned layman ($N = 4$), did not provide enough data, and the covariance matrix had no inverse. Therefore, no Mahalanobis distance could be calculated in the main study, and we had to drop this profile in the analysis.



Figure 3: Possible PPA implementations (a) Notification PPA, (b) Recommendation PPA, and (c) Decision PPA.

4.1.3 Privacy Profile Assignment. In this part of the survey, participants answered the questionnaire for the assignment of a privacy profile. For this, we used the questionnaire from our pre-study. After that, we directly proceeded to the next study part without showing the profile description without the profile titles to the participants.

4.1.4 PPA Implementation Possibilities General. After the profile assignment, we introduced the general concept of a PPA to the participants.

To do so, we built upon the qualitative study by Colnago *et al.* [13] which specifically investigated three different implementations for PPAs for IoT in general. Hence, we introduced three types of PPAs to the participant: 1) notification PPA, 2) recommendation PPA, and 3) decision PPA. We adapted the descriptions from Colnago *et al.* [13] and provided the participants with screenshots of possible PPA implementations on a mobile phone to support the understanding of the PPA (see Figure 3). After the PPA descriptions, we asked nine control questions to check whether the participants understood how the PPA functions. Then, we asked the participants to rank the three PPAs according to their preferences and to explain their ranking in a free-text field. In this part, we first wanted to focus on general perceptions of the different PPA implementations before introducing different scenarios.

4.1.5 PPA Implementation Possibilities in Scenarios. In this part, we introduced different scenarios of a PPA that varied in three aspects:

- (1) the environment in which the IoT device is located (environment),
- (2) the captured data (data), and
- (3) the number of data requests (request).

The factors were chosen through a focus group discussion with seven experts in the field of IoT privacy. The experts discussed and chose factor variations based on related work and the study constraints which were given by the study duration, test economy, and evaluability of the results. Table 1 lists all factor variations. The number of requests was informed by the scalability aspects from [13, 47], and the specific numbers were chosen from [27] to reflect

a spectrum of the number of notifications a user might receive throughout a day. The environments and data types were chosen based on the results from previous studies [6, 19, 38, 54]. As a result, we investigated two types of environments, three levels of data sensitivity, and three levels of privacy decisions. Combining these three factors leads to a total of 18 possible combinations of contexts. Each participant received all 18 combinations in random order.

Each scenario was presented to the participant as follows. First, they received a description of the scenario in the form (Table 4 in Appendix B lists all scenario statements that were presented to the participants):

Imagine you are at [environment] and your own IoT devices request access to [data]. This request happens [request] a day (this means about [number] times per hour).



Then, we asked the participants to rank the three PPAs and whether they would like to allow or deny the data collection in this scenario.

4.1.6 Demographics. Finally, we asked for demographics.

4.2 Recruitment and Participants

The sample consisted of 1126 participants residing in different countries. For recruitment and reimbursement, we used the online platform Prolific. 423 of them identified as female, 690 as male, seven as other, and six preferred not to say. The participants were between 18 and 68 years old ($Mean = 26.88$, $SD = 9.12$). 40.05% of participants were full or part-time employees, 36.69% were students, 19.44% were unemployed, or retired, and 0.79% preferred not to say. The overwhelming majority of 1116 participants reported daily Internet usage. The remaining participants stated values between 4 to 6 times a week and less than once a week. 41.74% ($N = 470$) participants previously heard of the Internet of Things before the study, while 58.26% ($N = 646$) have not.

Table 1: Overview of the factors that we varied in the investigated scenarios.

Factors	Levels	Explanation
environment	at home  , at another household 	location where data is captured
data	biometric data  , consumer behavior data  , presence 	data that is captured
request	low (5× per day), medium (25× per day), high (100× per day)	the number of requests for data capture

4.3 Ethical Considerations

All studies reported in this paper were conducted in line with the recommendations from the ethics committee at our institution. Before each study, participants were shown a consent form containing information about the study’s goal, its procedure, and the study’s data protection policy. Data collection and storage were in line with national and regional data protection laws in our country. Furthermore, we provided the participants with contact information from the examiners and researchers. The participants were compensated based on a \$12/hour rate.

4.4 Methodology Limitations

In this section, we explain the limitations of our study. The Dupree profiles were assigned based on qualitative data [17]. Since our goal was to develop an assignment method that could be realistically implemented in a PPA, we decided to use a quantitative assignment method using a limited number of items. Of course, such a method cannot be a substitute for rich qualitative data. However, assigning a privacy profile based on qualitative data would not have matched the research goal of developing an assignment questionnaire. Further, similar to [39], we aimed to group users in the very narrow scope of (not) granting access to data requested by IoT devices in private environments (at home, a friend’s home, an AirBnB). While this can be used to reduce burden from users, it might also overly simplify the complex nature of privacy decisions, which include more context than the location, collected data and number of requests. Based on our study design, we presented all combinations of factors to each participant resulting in 18 conditions. While the factors investigated by us form important information for making privacy decisions, there are further factors not investigated in our study that might have an impact, such as the provider of the IoT device [82], or the owner of the IoT environment [45]. Consequently, more research needs to be done before putting PPAs into action.

Even though we used a recruitment platform, our sample is disproportionately male and young, with a high percentage of students. This leaves open the question of what the distribution of personas and preferences for a PPA would have looked like in a balanced sample. However, the similarity of distribution within the studies indicates that the sample in the main study was not impacted. In our studies, self-reported answers to questions were collected. These answers may not reflect participants’ preferences in the real world. We addressed this by using items from well-established constructs. However, future research should verify our results with additional data sources.

5 MAIN STUDY RESULTS

We collected a data set with 1228 complete records. Then, we checked the control items and removed 102 participants who answered at least two of the nine control items incorrectly. We used

nine items distributed over the entire questionnaire to ensure data quality. The first type of item was ensuring participants understood the PPA concept correctly by asking them six questions about the different PPA implementations. The second set of items were three attention checks asking participants to click on a specific answer. The final data set consists of 1126 complete responses. Participants took an average of 18.12 minutes to complete the survey.

5.1 Distribution of Privacy Profiles

We assigned the profile by calculating the vector distances using the Mahalanobis distance and the covariance matrix calculated from the pre-study III data. Figure 4a provides an overview of the distribution of privacy profiles.

73.89% of profile assignments belong to a privacy profile from Dupree *et al.*, while 294 assignments (26.11%) correspond to profiles from our extension. The struggling amateur from Dupree *et al.* was the most prominent privacy profile (35.34%, $N = 398$), followed by the technician (17.94%, $N = 202$). Another prominent privacy profile was the lazy amateur from our extension (17.14%, $N = 193$). Fifteen (1.33%) assignments corresponded to the motivated layman.

5.2 General PPA Preferences

We asked the participants to rank the three PPA implementations based on their preferences and to explain the ranking. Figure 4b shows the distribution of the top choice.

5.2.1 Overall: Independent from Privacy Profiles – Quantitative Results. Decision PPAs were most frequently placed in the last position ($N = 648$, 57.54%). Recommendation PPAs were most frequently placed in the second position ($N = 516$, 45.78%) and slightly less often placed in the top one ($N = 465$, 41.29%). The notification PPA was selected similarly often on the first ($N = 399$, 35.43%) and second ($N = 394$, 34.99%) position. We analyzed this choice with a Pearson χ^2 -test that revealed significant differences ($\chi^2(2)=57.1$, $p<0.001$). For the post-hoc analysis, we calculated pair-wise χ^2 -tests and applied Bonferroni correction (corrected α of .016). The differences between notification and decision assistants and those between recommendation and decision assistants were significant (each $p < .001$).

5.2.2 Overall: Independent from Privacy Profiles – Qualitative Results. When asked to justify their decisions, participants named arguments for and against each PPA implementation.

Data Analysis: We analyzed the free-text answers using inductive qualitative content analysis according to Kuckartz [32]. In this analysis approach, the qualitative content is structured through coding text entities into thematic categories and related subcategories, letting patterns and categories emerge naturally from the data. To ensure adequate reliability, the formation and assignment




Figure 4: (a) Overview about the distribution of privacy profiles in our sample ($N = 1126$), (b) overall preferences of PPA implementations. The asterisk * indicates statistically significant differences.

of categories were discussed with two other researchers in review meetings during the evaluation. Data segments were constantly compared, and final code allocations were agreed on among the researchers. We report our findings by presenting themes and categories that have emerged from the data.

Decision PPAs: Participants favoring decision PPAs stated *convenience-related aspects* as users must *only decide once* ($N = 36$). It also *decides automatically* ($N = 45$), so users *do not have to do anything* ($N = 23$), *saving time* ($N = 31$). P193 said “*I think the decision assistant software is the best option because I spend the least time interacting with it.*”

However, there are also many arguments against decision PPAs. Participants *want to be informed* ($N = 58$) about 1) which data is collected, 2) which devices are in the room, and 3) which decisions are made by the PPA. Participants wanted to *make conscious decisions* ($N = 211$) and *maintain control* ($N = 43$). Another aspect was *missing trust towards AI* ($N = 51$). Also, privacy-related aspects were mentioned. Participants do not want data about them to be collected or processed by the PPA ($N = 34$). Sample comments against using decision PPAs are P204: “*it loses all the idea of privacy by taking the choice and control from the user.*” or P214: “*I really don’t want to give up all decision-making process for AI.*”

Notification PPAs: An advantage of notification PPAs is that participants are *informed* ($N = 124$), e.g., about present IoT devices. At the same time, they *retain control* and can make their own decisions ($N = 203$): “*Because it allows you complete freedom regarding the choices you could make.*” (P368).

However, participants noted that notification PPAs offer *fewer possibilities and information* compared to the other PPAs ($N = 59$), while *annoying* users with notifications or demanding decisions ($N = 75$). The *support* offered by notification PPAs is insufficient ($N = 34$). On the other hand, it *demands more attention and time* from the user ($N = 48$).

P529 wrote “*I reject the Notification Assistant because I would need to read very carefully the notification every time I receive it to be coherent with my preferences.*” while P624 commented “*And the notification, like I’ve said before, wouldn’t be helpful if you don’t know what you’re doing.*”

Recommendation PPAs: Participants liked the recommendation PPA because it *makes suggestions* based on previous choices or preferences ($N = 104$) while *providing full control* ($N = 258$). This was especially helpful for persons who lack knowledge or experience in privacy protection ($N = 22$). P601 stated “*because it informs me about the devices and reminds me of my preferences, but still lets me choose for myself.*”

Similarly to decision PPAs, participants received it as negative that recommendation PPAs might *store and process data* about users’ preferences ($N = 20$). Also, some participants worry that the recommendations might be *biased* ($N = 14$). For some participants, the notifications by recommendation PPAs are perceived *unnecessarily long* ($N = 12$), e.g., P413: “*The recommendation assistant is collecting your data of preferences, like the decision assistant software.*” or P588: “*I prefer selecting my preferences myself than using what the device thinks I might like.*”

Result Summary: In summary, we can conclude that, overall, recommendation PPAs were preferred by the participants when the analysis is not linked to the specific privacy profile or scenario. Yet, participants also named different reasons for (dis)liking all of the presented PPAs.

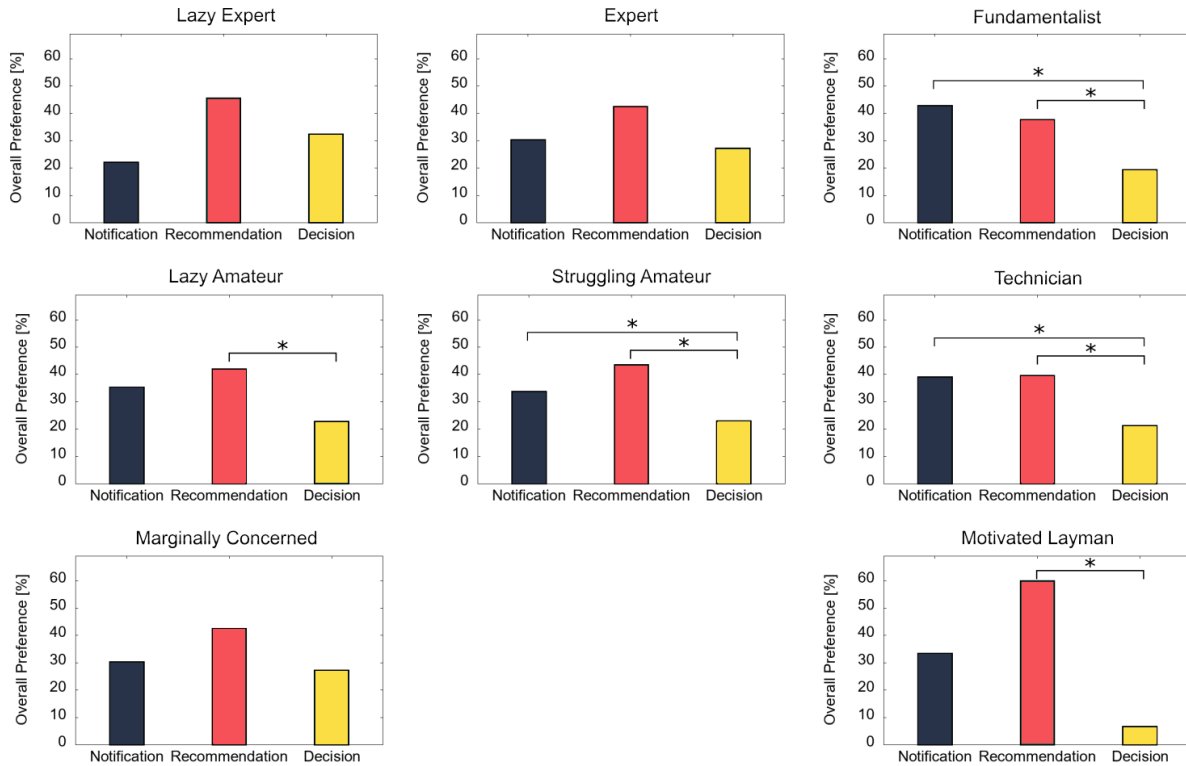
5.2.3 Profile-Specific: Dependent from Privacy Profiles – Quantitative Results. In this section, we consider which PPAs are preferred based on the assigned privacy profiles. We analyzed each profile cluster by a Pearson χ^2 -test. If this test revealed significant differences, we proceeded by calculating pairwise χ^2 -tests in the post-hoc analysis applying Bonferroni correction (corrected significance level .016). Considering the Pearson χ^2 -tests the differences between the PPA choices in the profiles lazy expert and expert were not significant, while all others were significant (see Table 2). The results of the post-hoc analysis are provided in Fig. 5.

Overall, the results resemble the choices that were not linked to privacy profiles with a tendency to recommendation PPAs. Considering fundamentalists, recommendation and notification PPAs were significantly preferred over decision PPAs. A majority of participants in each profile did not prefer decision PPAs.

5.2.4 Profile-Specific: Dependent from Privacy Profiles – Qualitative Results. We analysed the qualitative data as detailed in Sec. 5.2.2.

Table 2: Overview of the χ^2 -tests conducted on the main study findings to analyse differences in the preferred PPA based on the assigned privacy profiles.

	Fundamentalist	Technician	Motivated man	Lay-	Lazy Expert	Lazy Amateur	Marginally concerned	Expert	Struggling Amateur
χ^2	8.96	13.2	6.4		5.68	11.0	23.1	2.88	25.4
df	2	2	2		2	2	3	2	2
p	.011	<.001	.041		.059	0.004	<.001	.236	<.001

**Figure 5: Overview of PPA preferences based on the assigned privacy profile. Asterisks * indicate statistically significant differences.**

Fundamentalists Prefer Notifications: Fundamentalists justified preferring notification PPAs since they offer the best control and do not require data from its users. e.g., P1403 wrote “Being an advanced user, I prefer to not have my mobile device make any recommendations or make any decisions which it thinks is best for me. I prefer to make the decisions my self and not rely on my device. I chose the notification assistant first as I would like to be notified every time I enter a room and choose what I prefer; sometimes I allow the voice to be recorded and sometimes I would not.” while P257 commented “I prefer the notification assistant so that I can choose for each situation independently. Is does not matter what my previous choices were, I prefer to analyze each of them separately. In this sense, the decision assistant software would not be a good option. Also I think this kind of software makes people ‘lazy’ about their decisions.”

Low Knowledge Profile Value Recommendations: Participants with a profile that represents low knowledge valued the support by

the recommendation PPAs to make privacy-friendlier decisions while offering control; such as P259: “Ideally, I would like complete autonomy when deciding the extent to which data is collected from me, however, I am a beginner when it comes to protecting my data. Therefore, I believe a Recommendation Assistant would helpful since it can suggest options to me while still giving me the final choice.” or P236: “My one big worry with the Decision Assistant Software would be the effectiveness and how it ranks decisions. I would be worried about compromises given to malicious devices or completely new devices that perhaps the software does not recognize or is somehow taking advantage of the software in some way. I would want a choice rather than leaving it up to a software’s “better judgement.”

Result Summary: Overall decision PPAs were least preferred by all profiles. Fundamentalists and technicians would even prefer notification PPAs based on privacy preferences as recommendation PPAs require user data.

5.3 PPA Preferences based on Usage Contexts

In this section, we present the results of the 18 scenarios. To evaluate the PPA choice in these scenarios, we first analyzed our data independently from the privacy profile assignment to obtain *overall results*. Then, we analyzed the scenarios *based on the assigned privacy profiles*.

5.3.1 Overall: Independent from Privacy Profiles. First, we analyzed the impact of the scenarios on the choice of PPA implementations considering the data, environment, and request levels are independent variables. For this, we analyzed our data with a Friedman test revealing significant differences for data ($\chi^2(8)=560.48$, $p<.001$), environment ($\chi^2(5)=370.85$, $p<.001$), and number of requests ($\chi^2(8)=949.86$, $p<.001$). In the next step, we calculated pairwise comparisons and applied Bonferroni correction (see Table 3 in Appendix B). All comparisons were significant except for decision and recommendation PPAs considering a medium (25) number of requests, home, video, audio, biometric data, and consumer behavior. In the remainder of this section, we report the results based on the individual factors.

Request: For a low (5) number of decisions, the participants placed all PPA implementations equally often on top of the ranking. In the medium level, 49.39% placed decision PPAs on top, while 61.42% placed decision PPAs on top in the high level. In cases with high requests, decision PPAs were preferred over recommendation PPAs ($z = 10.16$, $p < .0001$). If the number of requests is low, recommendation PPAs were preferred over notification PPAs ($z = 7.595$, $p < .0001$). Based on these results, we conclude that the number of requests is essential for choosing PPAs.

Environment: In scenarios with other households, participants generally preferred recommendation PPAs (45.43%) significantly more often than decision PPAs ($z = -3.60$, $p = .005$).

Data: Considering the captured data, we could not find statistically significant differences.

Data Collection Preferences: In each scenario, we asked participants whether they would allow or deny data collection. Over all scenarios, 41.31% of data collections were allowed, and 68.69% were denied. Considering the collected data, 24.75% allowed the collection of biometric data, 45.31% allowed the collection of consumer behavior, and 45.66% allowed to collect presence data. At home, 46.54% of data collection was allowed. In contrast, 36.06% of collections in other households were allowed. If data collection was requested frequently, participants denied more data collection. The low (5) number of requests resulted in 45.66% of allowed collections, medium (25) in 41.65%, and high (100) in 36.62%.

5.3.2 Profile-Specific: Dependent from Privacy Profiles – Implementation Choice. After analyzing overall preferences, we investigated the relations between the privacy profiles and the PPA choices by analyzing the rankings. Figure 6 lists the top PPAs choices based on privacy profile and scenario. None of the privacy profiles exclusively chose a single PPA implementation over all scenarios. The profile marginally concerned (low motivation and knowledge) almost exclusively preferred decision PPAs except for scenarios in other households with about five requests per day (S10 and S16).

In the scenarios S2, S3, S8, S9, and S18 all privacy profiles preferred decisions PPAs. All of these scenarios were at home with more than 25 requests per day. The collected data were video, audio, and biometric (S1 and S2), and consumer behavior (S8 and S9). For all other at-home scenarios, the profiles also preferred decision PPAs, except for the profile ‘motivated layman’. Motivated laymen tended to favor notification and recommendation PPAs if the number of requests is low (5) daily. However, the share of the motivated layman in our study was quite low (Fig. 4a), and the second and third-ranking was almost equally often.

Experts and fundamentalists (both highly knowledgeable and at least medium motivated) preferred notification PPAs if the number of requests is low (5) (S1, S4, S7, S10, S13, S16). However, looking at the second-ranking, experts, that are less motivated than fundamentalists, also chose decision PPAs for consumer behavior data at home (S7). Lazy experts (high knowledge, low motivation) preferred decision PPAs in all scenarios except for those at another household with video, audio, and biometric data (S4) and presence data (S16). In those scenarios, they prefer recommendation PPAs.

Result Summary: Overall, the number of requests mostly influenced the participants’ choices of PPAs. Considering the privacy profiles, if the motivation is higher recommendation PPAs are chosen. At the same time, if the knowledge is higher, there is a tendency to choose notification PPAs.

5.3.3 Profile-Specific: Dependent from Privacy Profiles – Data Collection Choice. In each scenario, we asked the participants whether they would allow or deny the data collection. Figure 7 provides an overview considering whether the majority of users within the privacy profile allowed or denied the collection. Figure 8 extends this information by depicting the allowance rate. Each privacy profile demonstrated an individual pattern for allowance and denial indicating that the profiles differ from each other.

In the scenarios S1–S6, the collected data was video, audio, and biometric. All privacy profiles denied data collection. In these scenarios, the allowance rates were generally low. The rates in other households were lower compared to the own home. Privacy profiles with a high motivation (fundamentalist, technician, motivated layman) denied most data collections except for S13, S14, and S16. These scenarios share a low (5) number of requests.

S13 and S14 were allowed by each privacy profile and correspond to presence sensor data at home. The profile marginally concerned placed the least restrictions on data collection by allowing everything except for S1–S6, S12, and S18. The latter two refer to a high (100) number of requests in foreign households.

Result Summary: For making the decision whether to allow or deny a data collection request, the participants mostly considered the collected data (i.e., biometric, consumer behavior, presence). This is followed by the environment since data collection at home was more frequently allowed compared to other households. Considering the privacy profiles, the differences were mainly rooted in the motivation. The fewer motivation and knowledge a privacy profile represents, the more data collections are allowed.




Figure 6: Ranking results grouped by privacy profiles and scenarios. Each profile has an individual pattern. In scenarios with at least medium request, decision PPAs were favored. In other scenarios, profiles representing high knowledge and motivation tend to prefer notification PPAs. Lower motivation led to the choice of recommendation PPAs. Legend: at home; foreign household; video, audio & biometric data; consumer behavior; location data; the numbers denote the number of requests per day.




Figure 7: Decisions of privacy profiles to allow or deny data collections in the usage contexts based on the majority. Each privacy profile demonstrated an individual pattern of allowance and denial. Differences are mainly connected to motivation. Legend: at home; foreign household; video, audio & biometric data; consumer behavior; location data; the numbers denote the number of requests per day.




Figure 8: Allowance rate of data collections based on the privacy profiles and the usage contexts. Each privacy profile has an individual pattern of allowance rates. The allowance rate in S1 to S6 represent the collection of video, audio, and biometric data. However, the allowance rate in other households is lower than at home. Differences are mainly connected to motivation. Legend: at home; foreign household; video, audio & biometric data; consumer behavior; location data; the numbers denote the number of requests per day.

5.4 Further Factors for Making Privacy Decisions

After presenting all scenarios, we asked participants whether the presented information was sufficient to make privacy decisions and whether they missed any information. Most participants (65.36%) stated that the provided information was sufficient. The remaining participants named the following aspects. Some participants (6.5%) were interested in information about the *purpose* of the data collection, *who* can access the data, *who* collects the data, *where* it is stored, *how long* it is stored. Other participants (2.1%) did not miss information but asked about the consequences of not reacting to a notification. For instance, P735 said *"I would have liked to know the default action when a notification is not answered. If the notifications assistants automatically denied access to ignored requests, they might have been more desirable."* Others had suggestions regarding PPA implementations or provided an answer unrelated to the question.

6 DISCUSSION

In this section, we first discuss the results with regard to our *main research questions* to provide recommendations for implementing PPAs. Next, we discuss clustering in general. Finally, we conclude with *guidance* for future studies.

6.1 RQ1: How can privacy profiles be used to create autonomous PPAs?

To answer RQ1, we performed three pre-studies, totaling 417 participants. We built upon previous work on privacy profiles by Dupree *et al.* [17] and realized that the profiles do not cover all levels of knowledge and motivation. Based on our extension of the Dupree personas, we developed a short assignment questionnaire and validated it in the third pre-study showing that participants in general considered the proposed privacy profiles to fit their privacy needs. Hence, we conclude that the short questionnaire can serve as a basis for assigning a privacy profile. While we specifically considered privacy profiles, other options exist for realizing PPA decisions. Machine learning (ML) approaches in the literature used past decisions to predict new ones [19] (unsupervised ML) or segment large data sets of decisions into clusters [41] (supervised ML). In this context, our work is closest to unsupervised ML.

The approach based on past decisions requires data from the users and might not be predictable enough to affect user trust. While decision and recommendation PPAs can have the same drawbacks regarding privacy, notification PPAs do not require any knowledge about their users to function. Using data from other individuals might result in a misclassification of the user. Decision PPAs can be more predictable than ML-based approaches depending on the decision model. Similar to the misclassification by ML, a wrong privacy profile might be assigned. As can be evidenced by the information above, there are benefits and drawbacks to using ML or privacy profiles. ML might even be part of an approach based on profiles.

6.2 RQ2: Which PPA decision support types are preferred?

Most participants in the main study generally preferred recommendation PPAs informing them about data collection in their

environments and recommending a decision. Thus, participants stay in control and aware of the data collection. These results confirm the qualitative study by Colnago *et al.* [13], who interviewed 17 participants about PPA implementations. When looking at the scenarios in more detail, the choice for a particular PPA is more nuanced. Privacy decisions in specific situations can be different from overall preferences which confirms results from prior work [3, 19].

Privacy profiles representing users with high knowledge about privacy protection prefer notification PPAs if their motivation is at least medium. Such users would refrain from using recommendation PPAs because the data needed for giving a recommendation is considered privacy-sensitive. Furthermore, these users fear being manipulated by the recommended decision. Lazy experts form an exception. They are highly knowledgeable, but their motivation level is low. Hence, they prefer recommendations instead of notification PPAs. Users of all other privacy profiles generally preferred the recommendation PPA over the decision PPA especially because they want to be in control. For implementing PPA software, these results mean users should be kept in control by default. Either the behavior (notify, recommend, or decide) could be determined by the privacy profile, or users should be able to choose it. If no privacy profile is used, recommendation PPAs form a good standard or fallback setting that fits many users. Based on that, we formulate recommendation R1:

R1: If no context is provided, notification PPAs should be chosen for fundamentalists and experts. For all other profiles, recommendation PPAs should be provided.

6.3 RQ3: How do users' wishes for PPA support differ based on the context of use?

When we investigated different scenarios, we found that participants changed their PPA preferences. The frequency of requests impacted users' decisions the most. We also found different tendencies in the scenarios connected to the level of knowledge and motivation represented by the profile. Within this scope, our results indicate that users' motivation is more important than their knowledge.

6.3.1 General Observations. Participants primarily considered the number of requests when deciding on a PPA. If the number of requests per day is medium (25) or high (100), participants opted for a decision PPA that can take over work in frequent privacy decisions. Hence, users are more likely to trade control for convenience by preferring decision PPAs. Consequently, scalability was identified as the most important factor in our investigation.

In the remainder, we discuss the impact of the privacy profiles on PPA choice based on the knowledge about privacy protection represented by the profiles. For implementing a PPA, this means that the number of data collection requests should be kept low to keep users in control. Ideally, the number of requests by the device would be either minimized or the user should be offered a "remember me" function, storing the answer for future decisions. Thus, the user's decision is captured by a notification or recommendation assistant but automatically communicated in the future. We suggest:

R2: The number of privacy decisions and notifications should be as low as possible to keep users in control.

R3: Users should be offered a "remember me" function for privacy decisions that they can activate if they wish.

6.3.2 Profiles with High Knowledge. Profiles with high knowledge and at least medium motivation prefer notification PPAs. These PPA implementations demand the highest user involvement. In scenarios with at least a medium (25) number of decision requests, users of those profiles shifted to decision PPAs as convenience became more important. The relation of convenience and privacy was investigated by prior work showing that convenience is a major factor for sacrificing privacy [19, 82]. However, the sacrifice of privacy is not towards the IoT device but towards the PPA software. Also, in this scope, lazy experts form an exception based on their motivation. They tend to prefer decision PPAs in each scenario and shift to recommendation PPAs in other households with few (5) decision requests.

6.3.3 Profiles with Medium Knowledge. Technicians, struggling amateurs, and lazy amateurs are profiles with medium knowledge. Both amateur profiles (low and medium motivation) generally preferred recommendation PPAs in scenarios with low (5) requests. At home, they even tended to decision PPAs. In contrast to these results, technicians with high motivation, tended towards notification PPAs in scenarios with few requests. For consumer behavior in other households, they prefer to receive a recommendation. Considering presence data at home, they opt for notification PPAs. This matches with previous work on bystander privacy in which users underestimated the impact of data collection in other households [45]. We conclude that the support should be tailored to the location of the user offering more control in private environments. This leads to our next recommendations:

R4: The PPAs of users with high privacy knowledge and at least medium motivation should be designed as follows: If the users can manage the number of requests, the PPAs should not collect data (notification PPAs). If the number of requests exceeds the manageable level, the PPAs can make decisions for their users (decision PPAs). Ideally, users are involved in making the decision when to transition from notification to decision PPA.

R5: The communication and support offered by the PPA should be tailored to the user's location. If the location is their own home, users should be given the most control.

6.3.4 Profiles with Low Knowledge. The last profiles are those with low knowledge, namely, marginally concerned and motivated laymen. Marginally concerned users almost exclusively chose decision PPAs in other households with low (5) requests per day (S10 and S16). In contrast, motivated layman differentiated between their own homes in which they tended to prefer notification PPAs except for access to presence data, and other households in which they either want a notification or recommendation if the number of requests is medium (25) or low (5). This also support our recommendation R5. The knowledge of users is crucial in the choice of a PPA. While the most support is given by the decision PPA, what the PPA does might not be apparent to the user. At the same time, notification PPAs might overwhelm the users. We argue that

recommendation PPAs could be leveraged to educate users about privacy decisions and their consequences. Instead of basing the recommendation on past user decisions, the recommendations could be given by data from privacy experts. Based on that, we provide the following recommendations:

R6: Recommendation PPAs can be used to educate users with low privacy knowledge about privacy-friendly decisions while keeping them in control. If the users have low motivation, decision PPAs should be provided.

6.4 RQ3: How do decisions for allowing/denying data collection differ among different privacy profiles?

Each privacy profile resulted in an individual pattern representing the allowance and denial of data collection in IoT scenarios. Similar to a study by Naeini *et al.* [19], participants generally differentiated between their private homes and other private households as reflected in R5. This differentiation was common among all profiles. Furthermore, access to presence data at home was allowed by each profile if the number of requests was about 25 per day. Higher requests were denied by privacy profiles that represent high motivation. Those were also the profiles that denied most data collection. While we found individual patterns, our data has to be seen as an approximation because more factors can be considered while making privacy decisions. Further, even if the patterns were individual, we could observe the common trend to deny the collection of audio, video, and biometric data. Here, a more nuanced investigation is required to draw conclusions, although we can observe a connection to motivation. We provide the following recommendation:

R7: A PPA should be based on a rich set of decision factors, such as location, device manufacturer, data type, and so on. Users should be able to choose the factors that they want to be informed about.

R8: The own home is considered as more privacy-sensitive compared to other households. Hence, users should be offered the highest degree of control at home (no decision PPAs).

6.5 Clustering versus Individual Solutions

The general goal of profiles or personas is modeling the needs and attitudes of user archetypes that represent a group of users with similar needs. This results in a simplification of reality to make complex processes manageable. However, the simplification also has drawbacks, as it might create stereotypes that overly simplify individual differences. As privacy in itself is a highly individual concept [15] deeply rooted in an individual's personal ideals, needs, and choices, profiles create tension. This results in the question of whether profiling should be used when implementing PPAs. Current privacy settings in most software solutions, e.g., browsers or smartphones, tend to overwhelm users even if they are guided through them in a setup procedure [30]. Because of that, we risk that a majority of individuals use one single set of settings, neglecting individual needs. Considering this starting point, privacy profiles offer more individual solutions even if they represent groups of users and are applied to a very narrow scope [39]. Further, profiles

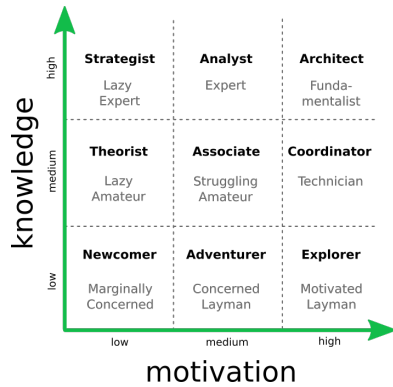


Figure 9: Proposed persona titles: Our proposed titles are shown in bold letters. The current titles are shown below the proposed ones.

might be a good starting point for users to explore different kinds of settings. In sum, while we acknowledge that profiles simplify reality, they also bear the potential to create more individual solutions if used the right way.

6.6 Guidance for Future Work

We discuss guidance for future work in the domains of privacy profiles and PPAs.

6.6.1 Decision Factor Model. As for usage scenarios, we investigated different environments, data types, and number of decision requests based on information provided in related work [6, 13, 19, 27, 38, 47, 54]. Based on our study design, we presented all combinations of factors to each participant resulting in 18 conditions. While the factors investigated form important information for making privacy decisions, there are further factors not investigated in our study that might have an impact, such as the provider of the IoT device [82], or the owner of the IoT environment [45]. These factors could be investigated in the future to inform a model of privacy decision-making for each privacy profile.

6.6.2 Profile Communication. PPAs based on privacy profiles can be realized in two ways. First, the support can be based on the mobile device from the user that knows the user’s privacy profile. Second, the user could use privacy profiles more explicitly to communicate privacy preferences. This could be done by using a control device provided in the environment or even by tangible representations of the profile visually placed in specific spots. The mobile device does not have to communicate with the IoT device and exchange information; the user acts more actively and hence has a better knowledge about what is going on. Future work should investigate these methods with different user types.

We acknowledge the work put into developing the first set of profiles by Dupree *et al.* [17]. However, the terminology used in the profile titles is overly judgmental about individuals, as we learned throughout our research project. We chose to omit the profile titles to treat our participants respectfully. This also shows that different profile titles are needed in the future to communicate the (extended)

profiles to users and participants. We propose to use more positive and inclusive language by renaming the profiles as shown in Figure 9.

6.6.3 Broadening Profile Assignment for Other Domains. We specifically investigated PPAs for private IoT-equipped environments. Our proposed assignment method for a privacy profile is not limited to that domain. Hence, it could be used for different purposes that require clustering users based on their privacy needs. Related works [59, 60] have already demonstrated that such clustering is useful in the scope of the profiles from Dupree *et al.* Therefore, future work could leverage our assignment questionnaire to use the privacy profiles for other domains, such as Internet privacy or privacy on mobile devices.

6.6.4 Notification Hierarchy. In our study, we investigated different numbers of decision requests. In our PPA concept, those would be represented by notifications on mobile devices. In their daily lives, users receive a plethora of different notifications [27, 54]. For instance, these can be messages from others, alarms, or system notifications. PPA notifications would form a new group because they affect the users’ privacy. Therefore, it forms an important task of future studies to investigate the hierarchy of received notifications and whether users consider privacy-related notifications as more or less important compared to others.

6.6.5 Control versus Burden. Several works that investigated privacy aspects of technology have revealed a privacy paradox [24]. This is related to the fact that users in general wish to protect their privacy. However, when confronted with specific usage contexts, they consider other aspects, such as convenience [79], as more important. Our work indicates that scalability is of importance in the scope of privacy decisions. Thus, future works and implementations of PPAs should consider the scalability aspects and specifically investigate how many decisions users are willing to make on their own before delegating decision-making.

6.6.6 Comparison to Other Approaches. We investigated persona-based profiles. However, related work demonstrated the development of profiles based on machine-learning methods [19, 39, 40] or experts. Future work should investigate methods to compare these approaches in terms of performance and accuracy but also based on human factors, such as adoption and user trust.

7 CONCLUSION

In this work, we presented an in-depth investigation of personalized privacy assistance (PPA) in different usage scenarios in private households. For this, we conducted four investigations, totaling 1543 participants. From our studies, we learned that in general, PPAs that notify the users about data collection and those who recommend a decision whether to allow or deny this collection were favored. The reason for that was that users wanted to keep control. Users with profiles with high knowledge about privacy protection and at least medium motivation would even consider the data collected by PPAs to make the recommendation as either privacy-invasive or manipulating and thus prefer only to receive notifications. In contrast to that, considering the different usage contexts, users tended to choose decision PPAs that automatically make

and communicate decisions for them of the number of decision requests is at least 25 per day. This shows that convenience is valued more than remaining in control of every privacy-related decision. Our work serves as a stepping stone for providing personal privacy assistance based on privacy profiles; we describe the limitations of our approach and detail recommendations for implementing PPAs. Finally, while we specifically investigated PPAs in the scope of IoT environments, our proposed profile assignment questionnaire might be valuable for other domains since related works have demonstrated that the privacy profiles deliver value [59, 60].

ACKNOWLEDGMENTS

Parts of this work have been completed while Karola Marky and Verena Zimmermann were still affiliated with TU Darmstadt. This work was supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2092 CASA - 390781972 and grant number 251805230/GRK 2050, as well as the German Federal Ministry of Education and Research (BMBF) within the SWC 2.0 "PrivacyGate" 01S17050 and the Horst Görtz Foundation. This work was also supported by fif-Project Privacy Buddy and the National Research Center for Applied Cybersecurity ATHENE.

REFERENCES

- [1] Noura Abdi, Kopo M. Ramakapane, and Jose M. Such. 2019. More than Smart Speakers: Security and Privacy Perceptions of Smart Home Personal Assistants. In *Proc. of the Symposium on Usable Privacy and Security (SOUPS '19)*. USENIX Association, Berkeley, CA, USA, 1–16.
- [2] Tousif Ahmed, Roberto Hoyle, Patrick Shaffer, Kay Connelly, David Crandall, and Apu Kapadia. 2017. Understanding Physical Safety, Security, and Privacy Concerns of People with Visual Impairments. *IEEE Internet Computing* 21, 3 (May/June 2017), 56–63. <https://doi.org/10.1109/MIC.2017.77>
- [3] Marc Serramia Amoros, William Seymour, Natalia Criado, and Michael Luck. 2023. Predicting Privacy Preferences for Smart Devices as Norms. In *The 22nd International Conference on Autonomous Agents and Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS).
- [4] Noah Aporthe, Yan Shvartzshnaider, Arunesh Mathur, Dillon Reisman, and Nick Feamster. 2018. Discovering Smart Home Internet of Things Privacy Norms Using Contextual Integrity. *Proc. of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 2 (2018), 59. <https://doi.org/10.1145/3214262>
- [5] Gaurav Bansal, David Gefen, et al. 2010. The impact of personal dispositions on information sensitivity, privacy concern and trust in disclosing health information online. *Decision support systems* 49, 2 (2010), 138–150.
- [6] Debjanee Barua, Judy Kay, and Cécile Paris. 2013. Viewing and Controlling Personal Sensor Data: What Do Users Want?. In *Proc. of the International Conference on Persuasive Technology*. Springer, Cham, Switzerland, 15–26.
- [7] France Bélanger and Robert E Crossler. 2011. Privacy in the digital age: a review of information privacy research in information systems. *MIS quarterly* (2011), 1017–1041.
- [8] Ann-Renée Blais, Elke U Weber, et al. 2006. *Testing invariance in risk taking: A comparison between Anglophone and Francophone groups*. Technical Report. CIRANO.
- [9] Jürgen Bortz. 2013. *Statistik: Für Sozialwissenschaftler*. Springer-Verlag.
- [10] Denys Brand, Florence D. DiGennaro Reed, Mariah D. Morley, Tyler G. Erath, and Matthew D. Novak. 2019. A Survey Assessing Privacy Concerns of Smart-Home Services Provided to Individuals with Disabilities. *Behavior Analysis in Practice* (2019), 1–11. <https://doi.org/10.1007/s40617-018-00329-y>
- [11] Eun Kyoung Choe, Sunny Consolvo, Jaeyeon Jung, Beverly Harrison, Shwetak N. Patel, and Julie A. Kientz. 2012. Investigating Receptiveness to Sensing and Inference in the Home Using Sensor Proxies. In *Proc. of the Conference on Ubiquitous Computing (UbiComp '12)*. ACM, New York, NY, USA, 61–70. <https://doi.org/10.1145/2370216.2370226>
- [12] Hyunji Chung, Michaela Iorga, Jeffrey Voas, and Sangjin Lee. 2017. Alexa, Can I Trust You? *Computer* 50, 9 (2017), 100–104. <https://doi.org/10.1109/MC.2017.3571053>
- [13] Jessica Colnago, Yuanyuan Feng, Tharangini Palanivel, Sarah Pearman, Megan Ung, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. 2020. Informing the Design of a Personalized Privacy Assistant for the Internet of Things. In *Proc. of the 2020 CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3313831.3376389>
- [14] Lorrie Faith Cranor, Praveen Guduru, and Manjula Arjula. 2006. User Interfaces for Privacy Agents. *ACM Transactions on Computer-Human Interaction (TOCHI)* 13, 2 (June 2006), 135–178. <https://doi.org/10.1145/1165734.1165735>
- [15] Mary J. Culnan and Pamela K. Armstrong. 1999. Information Privacy Concerns, Procedural Fairness, and Impersonal Trust: An Empirical Investigation. *Organization Science* 10, 1 (1999), 104–115. <https://doi.org/10.1287/orsc.10.1.104>
- [16] Anupam Das, Martin Degeling, Xiaoyou Wang, Junjue Wang, Norman Sadeh, and Mahadev Satyanarayanan. 2017. Assisting Users in a World Full of Cameras: A Privacy-Aware Infrastructure for Computer Vision Applications. In *Proc. of the Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. IEEE, Piscataway, NJ, USA, 1387–1396. <https://doi.org/10.1109/CVPRW.2017.181>
- [17] Janna Lynn Dupree, Richard Devries, Daniel M. Berry, and Edward Lank. 2016. Privacy Personas: Clustering Users via Attitudes and Behaviors toward Security Practices. In *Proc. of the 2016 CHI Conference on Human Factors in Computing Systems (San Jose, California, USA) (CHI '16)*. ACM, New York, NY, USA, 5228–5239. <https://doi.org/10.1145/2858036.2858214>
- [18] Serge Egelman and Eyal Peer. 2015. Scaling the security wall: Developing a security behavior intentions scale (sebis). In *Proc. of the 33rd annual ACM conference on human factors in computing systems*. ACM, New York, NY, USA, 2873–2882.
- [19] Pardis Emami-Naeini, Sruti Bhagavatula, Hana Habib, Martin Degeling, Lujio Bauer, Lorrie Cranor, and Norman Sadeh. 2017. Privacy Expectations and Preferences in an IoT World. In *Proc. of the Symposium on Usable Privacy and Security (SOUPS '17)*. USENIX Association, 399–412.
- [20] Pardis Emami-Naeini, Henry Dixon, Yuvraj Agarwal, and Lorrie Faith Cranor. 2019. Exploring How Privacy and Security Factor into IoT Device Purchase Behavior. In *Proc. of the CHI Conference on Human Factors in Computing Systems (Glasgow, Scotland UK) (CHI '19)*. ACM, New York, NY, USA, Article 534, 12 pages. <https://doi.org/10.1145/3290605.3300764>
- [21] Cori Faklaris, Laura A Dabbish, and Jason I Hong. 2019. A self-report measure of end-user security attitudes (SA-6). In *Fifteenth Symposium on Usable Privacy and Security (SOUPS '19)*. Usenix, Berkeley, CA, USA.
- [22] Lujun Fang and Kristen LeFevre. 2010. Privacy Wizards for Social Networking Sites. In *Proc. of the 19th International Conference on World Wide Web (Raleigh, North Carolina, USA) (WWW '10)*. ACM, New York, NY, USA, 351–360. <https://doi.org/10.1145/1772690.1772727>
- [23] David Gefen. 2000. E-commerce: the role of familiarity and trust. *Omega* 28, 6 (2000), 725–737.
- [24] Nina Gerber, Paul Gerber, and Melanie Volkamer. 2018. Explaining the Privacy Paradox: A Systematic Review of Literature Investigating Privacy Attitude and Behavior. *Computers & Security* 77 (2018), 226–261. <https://doi.org/10.1016/j.cose.2018.04.002>
- [25] Nina Gerber, Benjamin Reinheimer, and Melanie Volkamer. 2018. Home Sweet Home? Investigating Users' Awareness of Smart Home Privacy Threats. In *Proc. of the Workshop on the Human Aspects of Smart Home Security and Privacy (WSSP '18)*. USENIX Association, Berkeley, CA, USA.
- [26] Yangyang He. 2019. Recommending privacy settings for IoT. In *Proc. of the 24th International Conference on Intelligent User Interfaces Companion - IUI '19*. ACM Press, Marina del Rey, California, 157–158. <https://doi.org/10.1145/3308557.3308732>
- [27] Shamsi T. Iqbal and Eric Horvitz. 2010. Notifications and Awareness: A Field Study of Alert Usage and Preferences. In *Proc. of the 2010 ACM Conference on Computer Supported Cooperative Work (Savannah, Georgia, USA) (CSCW '10)*. ACM, New York, NY, USA, 27–30. <https://doi.org/10.1145/1718918.1718926>
- [28] Michal Jakob, Zbynek Moler, Michal Pechoucek, and Roman Vaculin. 2011. Intelligent Content-Based Privacy Assistant for Facebook. In *Proc. of the IEEE/WIC/ACM International Conferences on Web Intelligence and Intelligent Agent Technology*. IEEE, Piscataway, NJ, USA, 499–500.
- [29] Timo Jakobi, Corinna Ogonowski, Nico Castelli, Gunnar Stevens, and Volker Wulf. 2017. The Catch(Es) with Smart Home: Experiences of a Living Lab Field Study. In *Proc. of the CHI Conference on Human Factors in Computing Systems (Denver, Colorado, USA) (CHI '17)*. ACM, New York, NY, USA, 1620–1633. <https://doi.org/10.1145/3025453.3025799>
- [30] Patrick Gage Kelley, Joanna Bresee, Lorrie Faith Cranor, and Robert W Reeder. 2009. A "nutrition label" for privacy. In *Proc. of the 5th Symposium on Usable Privacy and Security*. ACM, New York, NY, USA, 1–12.
- [31] Nadin Kökciyan, Pinar Yolun, et al. 2022. Taking situation-based privacy decisions: Privacy assistants working with humans. In *Proc. of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*. International Joint Conference on Artificial Intelligence, 703–709.
- [32] Udo Kuckartz. 2016. *Qualitative Inhaltsanalyse. Methoden, Praxis, Computerunterstützung (Grundlagentexte Methoden, 3., überarbeitete Auflage)*. Weinheim: Beltz Juventa. Zugriff am 9 (2016), 2017.
- [33] J. Sathish Kumar and Dhiren R. Patel. 2014. A Survey on Internet of Things: Security and Privacy Issues. *International Journal of Computer Applications* 90, 11 (2014).

- [34] Ponnurangam Kumaraguru and Lorrie Faith Cranor. 2005. *Privacy indexes: a survey of Westin's studies*. Carnegie Mellon University, School of Computer Science.
- [35] Hyosun Kwon, Joel E. Fischer, Martin Flinham, and James Colley. 2018. The Connected Shower: Studying Intimate Data in Everyday Life. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2, 4, Article 176 (Dec. 2018), 22 pages. <https://doi.org/10.1145/3287054>
- [36] Marc Langheinrich. 2002. A Privacy Awareness System for Ubiquitous Computing Environments. In *Proc. of the International Conference on Ubiquitous Computing*. Springer, Cham, Switzerland, 237–245.
- [37] Josephine Lau, Benjamin Zimmerman, and Florian Schaub. 2018. Alexa, Are You Listening?: Privacy Perceptions, Concerns and Privacy-Seeking Behaviors With Smart Speakers. *Proc. of the ACM Conference on Human-Computer Interaction 2, CSCW* (2018), 102. <https://doi.org/10.1145/3274371>
- [38] Scott Lederer, Jason I. Hong, Anind K Dey, and James A. Landay. 2004. Personal Privacy through Understanding and Action: Five Pitfalls for Designers. *Personal and Ubiquitous Computing* 8, 6 (2004), 440–454.
- [39] Jialiu Lin, Bin Liu, Norman Sadeh, and Jason I Hong. 2014. Modeling users' mobile app privacy preferences: Restoring usability in a sea of permission settings. In *10th Symposium On Usable Privacy and Security (SOUPS) 2014*. Usenix, Berkeley, CA, USA, 199–212.
- [40] Bin Liu, Mads Schaarup Andersen, Florian Schaub, Hazim Almuhamidi, Shikun Aerin Zhang, Norman Sadeh, Yuvraj Agarwal, and Alessandro Acquisti. 2016. Follow my recommendations: A personalized privacy assistant for mobile app permissions. In *Proc. of the Twelfth Symposium on Usable Privacy and Security (SOUPS)*. USENIX Association, Berkeley, CA, USA.
- [41] Bin Liu, Jialiu Lin, and Norman Sadeh. 2014. Reconciling Mobile App Privacy and Usability on Smartphones: Could User Privacy Profiles Help?. In *Proc. of the 23rd International Conference on World Wide Web (Seoul, Korea) (WWW '14)*. Association for Computing Machinery, New York, NY, USA, 201–212. <https://doi.org/10.1145/2566486.2568035>
- [42] Naresh K. Malhotra, Sung S. Kim, and James Agarwal. 2004. Internet Users' Information Privacy Concerns (IUPC): The Construct, the Scale, and a Causal Model. *Information Systems Research* 15, 4 (2004), 336–355. <https://doi.org/10.1287/isre.1040.0032>
- [43] Shrirang Mare, Franziska Roesner, and Tadayoshi Kohno. 2020. Smart Devices in Airbnbs: Considering Privacy and Security for both Guests and Hosts. *Proceedings on Privacy Enhancing Technologies* 2020, 2 (2020), 436–458.
- [44] Karola Marky, Nina Gerber, Michelle Gabriela Pelzer, Mohamed Khamis, and Max Mühlhäuser. 2022. "You offer privacy like you offer tea": Investigating mechanisms for improving guest privacy in IoT-equipped households. In *Proc. of the 22nd Privacy Enhancing Technologies Symposium (PETS)*. 400–420. <https://doi.org/10.56553/popets-2022-0115>
- [45] 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. Association for Computing Machinery, New York, NY, USA, 83–95. <https://doi.org/10.1145/3428361.3428464>
- [46] Karola Marky, Sarah Prange, Max Mühlhäuser, and Florian Alt. 2022. Roles Matter! Understanding Differences in the Privacy Mental Models of Smart Home Visitors and Residents. In *Proc. of the 20th International Conference on Mobile and Ubiquitous Multimedia (Leuven, Belgium) (MUM '21)*. Association for Computing Machinery, New York, NY, USA, 108–122. <https://doi.org/10.1145/3490632.3490664>
- [47] Karola Marky, Alexandra Voit, Alina Stöver, Kai Kunze, Svenja Schröder, and Max Mühlhäuser. 2020. "I don't know how to protect myself": Understanding Privacy Perceptions Resulting from the Presence of Bystanders in Smart Environments. In *Proc. of the NordiCHI Nordic conference on Human-computer Interaction (NordiCHI '20)*. ACM, New York, USA.
- [48] Karola Marky, Verena Zimmermann, Alina Stöver, Philipp Hoffmann, Kai Kunze, and Max Mühlhäuser. 2020. All in One! User Perceptions on Centralized IoT Privacy Settings. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI EA '20)*. ACM, New York, NY, USA, 1–8. <https://doi.org/10.1145/3334480.3383016>
- [49] Emily McReynolds, Sarah Hubbard, Timothy Lau, Aditya Saraf, Maya Cakmak, and Franziska Roesner. 2017. Toys That Listen: A Study of Parents, Children, and Internet-Connected Toys. In *Proc. of the 2017 CHI Conference on Human Factors in Computing Systems (Denver, Colorado, USA) (CHI '17)*. Association for Computing Machinery, New York, NY, USA, 5197–5207. <https://doi.org/10.1145/3025453.3025735>
- [50] Sarah Mennicken, David Kim, and Elaine May Huang. 2016. Integrating the Smart Home into the Digital Calendar. In *Proc. of the CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 5958–5969. <https://doi.org/10.1145/2858036.2858168>
- [51] Mateusz Mikusz, Steven Houben, Nigel Davies, Klaus Moessner, and Marc Langheinrich. 2018. Raising Awareness of IoT Sensor Deployments. In *Proc. of the Living in the Internet of Things: Cybersecurity of the IoT*. IET. <https://doi.org/10.1049/cp.2018.0009>
- [52] A. Morton and M. A. Sasse. 2014. Desperately Seeking Assurances: Segmenting Users by their Information-Seeking Preferences. In *Proc. of the Twelfth Annual International Conference on Privacy, Security and Trust*. IEEE, Piscataway, NJ, USA, 102–111.
- [53] Primal Pappachan, Martin Degeling, Roberto Yus, Anupam Das, Sruti Bhagavathula, William Melicher, Pardis Emami Naeini, Shikun Zhang, Lujo Bauer, Alfred Kobsa, Sharad Mehrotra, Norman Sadeh, and Nalini Venkatasubramanian. 2017. Towards Privacy-Aware Smart Buildings: Capturing, Communicating, and Enforcing Privacy Policies and Preferences. In *Proc. of the 37th International Conference on Distributed Computing Systems Workshops (ICDCSW)*. IEEE, Piscataway, NJ, USA, 193–198. <https://doi.org/10.1109/ICDCSW.2017.52>
- [54] Martin Pielot, Karen Church, and Rodrigo de Oliveira. 2014. An In-Situ Study of Mobile Phone Notifications. In *Proc. of the 16th International Conference on Human-Computer Interaction with Mobile Devices & Services (Toronto, ON, Canada) (MobileHCI '14)*. ACM, New York, NY, USA, 233–242. <https://doi.org/10.1145/2628363.2628364>
- [55] Frederic Raber, David Ziemann, Antonio Krueger, C. Weir, and M. Mazurek. 2018. The "Retailio" Privacy Wizard: Assisting Users with Privacy Settings for Intelligent Retail Stores. In *Proc. of the 3rd European Workshop on Usable Security (EuroUSEC '18)*. Internet Society, Reston, VA, USA.
- [56] Olivia K. Richards. 2019. Family-Centered Exploration of the Benefits and Burdens of Digital Home Assistants. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '19)*. ACM, New York, NY, USA, SRC11. <https://doi.org/10.1145/3290607.3308458>
- [57] Tom A. Rodden, Joel E. Fischer, Nadia Pantidi, Khaled Bouchour, and Stuart Moran. 2013. At Home with Agents: Exploring Attitudes Towards Future Smart Energy Infrastructures. In *Proc. of the SIGCHI Conference on Human Factors in Computing Systems (Paris, France) (CHI '13)*. ACM, New York, NY, USA, 1173–1182. <https://doi.org/10.1145/2470654.2466152>
- [58] Morris Rosenberg. 2015. *Society and the adolescent self-image*. Princeton University Press, Princeton, NJ, USA.
- [59] Manuel Rudolph, Svenja Polst, and Joerg Doerr. 2019. Enabling Users to Specify Correct Privacy Requirements. In *International Working Conference on Requirements Engineering: Foundation for Software Quality*. Springer, 39–54.
- [60] Manuel Polst Rudolph, Denis Feth, et al. 2019. Usable Specification of Security and Privacy Demands: Matching User Types to Specification Paradigms. *Mensch und Computer 2019-Workshopband* (2019).
- [61] Kathy S Schwaig, Albert H Segars, Varun Grover, and Kirk D Fiedler. 2013. A model of consumers' perceptions of the invasion of information privacy. *Information & Management* 50, 1 (2013), 1–12.
- [62] Susanne G. Scott and Reginald A. Bruce. 1995. Decision-Making Style: The Development and Assessment of a New Measure. *Educational and Psychological Measurement* 55, 5 (1995), 818–831.
- [63] William Seymour, Martin J. Kraemer, Reuben Binns, and Max Van Kleek. 2020. Informing the Design of Privacy-Empowering Tools for the Connected Home. *arXiv:2001.09077 [cs]* (Jan. 2020). <https://doi.org/10.1145/3313831.3376264> arXiv: 2001.09077
- [64] H. Jeff Smith, Sandra J. Milberg, and Sandra J. Burke. 1996. Information Privacy: Measuring Individuals' Concerns about Organizational Practices. *MIS quarterly* (1996), 167–196. <https://doi.org/10.2307/249477>
- [65] Yungpeng Song, Yun Huang, Zhongmin Cai, and Jason I. Hong. 2020. I'm All Eyes and Ears: Exploring Effective Locators for Privacy Awareness in IoT Scenarios. In *Proc. of the 2020 CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3313831.3376585>
- [66] Statista. 2019. Smart Home Worldwird. <https://www.statista.com/outlook/279/100/smart-home/worldwide> (Accessed January 2020).
- [67] Madiha Tabassum, Tomasz Kosinski, and Heather Richter Lipford. 2019. I Don't Own the Data": End User Perceptions of Smart Home Device Data Practices and Risks. In *Proc. of the Fifteenth Symposium on Usable Privacy and Security (SOUPS '19)*. Usenix, Berkeley, CA, USA.
- [68] Hua-Zhe Tan, Wei Zhao, and Hai-Hua Shen. 2018. A Context-Perceptual Privacy Protection Approach on Android Devices. In *Proc. of the International Conference on Communications (ICC)*. IEEE, Piscataway, NJ, USA, 1–7.
- [69] 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 *Proc. 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>
- [70] Blase Ur, Jaeyeon Jung, and Stuart Schechter. 2014. Intruders Versus Intrusiveness: Teens' and Parents' Perspectives on Home-entryway Surveillance. In *Proc. of the International Joint Conference on Pervasive and Ubiquitous Computing (Seattle, Washington) (UbiComp '14)*. ACM, New York, NY, USA, 129–139. <https://doi.org/10.1145/2632048.2632107>
- [71] Zixin Wang, Danny Yuxing Huang, and Yaxing Yao. 2023. Exploring Tenants' Preferences of Privacy Negotiation in Airbnb. In *Proc. of the 32nd USENIX Security Symposium (USENIX Security 23)*. USENIX Association, Anaheim, CA, 535–551.
- [72] Alan Westin. 1967. *Privacy and Freedom*. New York: Atheneum.

- [73] Lawrence R Wheelless. 1978. A follow-up study of the relationships among trust, disclosure, and interpersonal solidarity. *Human Communication Research* 4, 2 (1978), 143–157.
- [74] Maximiliane Windl, Albrecht Schmidt, and Sebastian S. Feger. 2023. Investigating Tangible Privacy-Preserving Mechanisms for Future Smart Homes. In *Proc. of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 70, 16 pages. <https://doi.org/10.1145/3544548.3581167>
- [75] Maximiliane Windl, Verena Winterhalter, Albrecht Schmidt, and Sven Mayer. 2023. Understanding and Mitigating Technology-Facilitated Privacy Violations in the Physical World. In *Proc. of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 585, 16 pages. <https://doi.org/10.1145/3544548.3580909>
- [76] Peter Worthy, Ben Matthews, and Stephen Viller. 2016. Trust Me: Doubts and Concerns Living With the Internet of Things. In *Proc. of the ACM Conference on Designing Interactive Systems (DIS '16)*. ACM, New York, NY, USA, 427–434. <https://doi.org/10.1145/2901790.2901890>
- [77] 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 *Proc. of the SIGCHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM, New York, NY, USA. <https://doi.org/10.1145/3290605.3300428>
- [78] Yaxing Yao, Justin Reed Basdeo, Oriana Rosata McDonough, and Yang Wang. 2019. Privacy Perceptions and Designs of Bystanders in Smart Homes. *Proc. of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–24.
- [79] Eric Zeng, Shirang Mare, and Franziska Roesner. 2017. End User Security & Privacy Concerns with Smart Homes. In *Proc. of the Symposium on Usable Privacy and Security (SOUPS '17)*. USENIX Association, Berkeley, CA, USA, 65–80.
- [80] Yu Zhai, Yan Liu, Minghao Yang, Feiyuan Long, and Johanna Virkki. 2014. A Survey Study of the Usefulness and Concerns About Smart Home Applications From the Human Perspective. *Open Journal of Social Sciences* 2, 11 (2014), 119. <https://doi.org/10.4236/jss.2014.211017>
- [81] Nicole Zhan, Stefan Sarkadi, and Jose Such. 2023. Privacy-enhanced Personal Assistants based on Dialogues and Case Similarity. In *European Conference on Artificial Intelligence*. IOS Press.
- [82] Serena Zheng, Noah Apthorpe, Marshini Chetty, and Nick Feamster. 2018. User Perceptions of Smart Home IoT Privacy. *Proc. of the ACM on Human-Computer Interaction* 2, CSCW (2018), 200. <https://doi.org/10.1145/3274469>
- [83] Verena Zimmermann, Merve Bennighof, Miriam Edel, Oliver Hofmann, Judith Jung, and Melina von Wick. 2018. 'Home, Smart Home'—Exploring End Users' Mental Models of Smart Homes. In *Mensch und Computer 2018-Workshopband*. Gesellschaft für Informatik e.V., Bonn, Germany, 407–417.
- [84] Verena Zimmermann, Paul Gerber, Karola Marky, Leon Böck, and Florian Kirchbuchner. 2019. Assessing Users' Privacy and Security Concerns of Smart Home Technologies. *i-com* 18, 3 (2019), 197–216. <https://doi.org/10.1515/icom-2019-0015>

- (1) I am highly motivated to protect my information privacy. (Privacy Motivation)
- (2) I know how to protect my information privacy. (Privacy Knowledge)
- (3) My own privacy is very important to me. (Privacy Importance)
- (4) I often am interested in articles about security threats. (Security Attitudes (SA-6))
- (5) I'm concerned that online companies are collecting too much personal information about me. (UIPC - Collection)
- (6) It is very important to me that I am aware and knowledgeable about how my personal information will be used. (UIPC Awareness)
- (7) I often procrastinate when it comes to making important decisions. (GDMS Avoidance)
- (8) I generally make important decisions at the last minute. (GDMS)
- (9) I make quick decisions. (GDMS Spontaneous)

A PRE-STUDY MATERIAL

In this section, we provide materials used within our pre-studies.

A.1 Evaluated Constructs in the Questionnaire Development

- (1) Westin's Privacy Segmentation Index (WPSI) [34]
- (2) Generalized Decision Making Style (GDMS) [62]
- (3) Rosenberg Self Esteem Scale (RSES) [58]
- (4) Internet Users Information Privacy Concerns (UIPC) [42]
- (5) Disposition to Trust [23]
- (6) Security Attitudes (SA-6) [21]
- (7) Consumer Alienation [61]
- (8) Domain Specific Risk Taking Skala (DOSPERT) [8]
- (9) Security Behavior Intentions Scale (SeBIS) [18]
- (10) Subjective Data Sensitivity [5]
- (11) Revised Self Disclosure Scale (RSDS) [73]
- (12) Privacy Motivation and Knowledge (self-formulated)

A.2 Assignment Questionnaire






























The questionnaire for assigning a privacy profile is as follows. The origin of the question is given in brackets.

B MAIN STUDY MATERIAL AND RESULTS

Table 3: Results of the post hoc tests for preferences of a PPA implementation in the different usage contexts. The adjusted significance is based on Bonferroni correction.

Factor	Pair	Test Statistic	z	Sig.	Adj. Sig.
high # requests	Decision Recommendation	1.172	10.157	0.000	0.000
	Decision Notification	2.761	23.919	0.000	0.000
	Recommendation Notification	1.588	13.762	0.000	0.000
medium # requests	Decision Recommendation	0.009	0.077	0.939	1.000
	Decision Notification	1.309	11.342	0.000	0.000
	Recommendation Notification	1.300	11.265	0.000	0.000
low # requests	Recommendation Notification	0.877	7.595	0.000	0.000
	Recommendation Decision	-1.425	-12.342	0.000	0.000
	Notification Decision	-0.548	-4.748	0.000	0.000
home	Decision Recommendation	0.120	1.526	0.127	1.000
	Decision Notification	1.095	13.883	0.000	0.000
	Recommendation Notification	0.974	12.356	0.000	0.000
other household	Recommendation Decision	-0.284	-3.604	0.000	0.005
	Recommendation Notification	0.857	10.875	0.000	0.000
	Decision Notification	0.573	7.271	0.000	0.000
presence	Decision Recommendation	0.167	1.447	0.148	1.000
	Decision Notification	1.468	12.719	0.000	0.000
	Recommendation Notification	1.301	11.273	0.000	0.000
video. audio. biometric	Recommendation Decision	-0.251	-2.174	0.030	1.000
	Recommendation Notification	1.309	11.342	0.000	0.000
	Decision Notification	1.058	9.168	0.000	0.000
consumer	Recommendation Decision	-0.080	-0.696	0.486	1.000
	Recommendation Notification	1.414	12.250	0.000	0.000
	Decision Notification	1.333	11.553	0.000	0.000

Table 4: Statements that describe the scenarios shown to the participants in the main study.

	Usage Context	Factor Levels
1	Imagine you are at home and your own IoT devices request access to video, audio, and biometric data about you (e.g., pictures of your face). This request happens five times a day (this means less than once per hour).	  5
2	Imagine you are at home and your own IoT devices request access to video, audio, and biometric data about you (e.g., pictures of your face). This request happens 25 times a day (this means 1-2 times per hour).	  25
3	Imagine you are at home and your own IoT devices request access to video, audio, and biometric data about you (e.g., pictures of your face). This request happens 100 times a day (this means 4-7 times per hour).	  100
4	Imagine you are in another household with IoT devices. The devices request access to video, audio, and biometric data about you (e.g., pictures of your face). This request happens five times a day (this means less than once per hour).	  5
5	Imagine you are in another household with IoT devices. The devices request access to video, audio, and biometric data about you (e.g., pictures of your face). This request happens 25 times a day (this means 1-2 times per hour).	  25
6	Imagine you are in another household with IoT devices. The devices request access to video, audio, and biometric data about you (e.g., pictures of your face). This request happens 100 times a day (this means 4-7 times per hour).	  100
7	Imagine you are at home and your own IoT devices request access to consumer behavior (e.g. whether a product has been consumed). This request happens five times a day (this means less than once per hour).	  5
8	Imagine you are at home and your own IoT devices request access to consumer behavior (e.g. whether a product has been consumed). This request happens 25 times a day (this means 1-2 times per hour).	  25
9	Imagine you are at home and your own IoT devices request access to consumer behavior (e.g. whether a product has been consumed). This request happens 100 times a day (this means 4-7 times per hour).	  100
10	Imagine you are in another household with IoT devices. The devices request access to consumer behavior (e.g. whether a product has been consumed). This request happens five times a day (this means less than once per hour).	  5
11	Imagine you are in another household with IoT devices. The devices request access to consumer behavior (e.g. whether a product has been consumed). This request happens 25 times a day (this means 1-2 times per hour).	  25
12	Imagine you are in another household with IoT devices. The devices request access to consumer behavior (e.g. whether a product has been consumed). This request happens 100 times a day (this means 4-7 times per hour).	  100
13	Imagine you are at home and your own IoT devices request access to data to a generic presence sensor that checks whether people are present in a room. This request happens five times a day (this means less than once per hour).	  5
14	Imagine you are at home and your own IoT devices request access to a generic presence sensor that checks whether people are present in a room. This request happens 25 times a day (this means 1-2 times per hour).	  25
15	Imagine you are at home and your own IoT devices request access to a generic presence sensor that checks whether people are present in a room. This request happens 100 times a day (this means 4-7 times per hour).	  100
16	Imagine you are in another household with IoT devices. The devices request access to a generic presence sensor that checks whether people are present in a room. This request happens five times a day (this means less than once per hour).	  5
17	Imagine you are in another household with IoT devices. The devices request access to a generic presence sensor that checks whether people are present in a room. This request happens 25 times a day (this means 1-2 times per hour).	  25
18	Imagine you are in another household with IoT devices. The devices request access to a generic presence sensor that checks whether people are present in a room. This request happens 100 times a day (this means 4-7 times per hour).	  100