

# Informing the Design of a Personalized Privacy Assistant for the Internet of Things

Jessica Colnago, Yuanyuan Feng, Tharangini Palanivel, Sarah Pearman, Megan Ung, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh  
Carnegie Mellon University, Pittsburgh, PA, USA  
{jcolnago, yuanyuanfeng, tpalanivel, spearman, mung, acquisti, lorrie, sadeh}@cmu.edu

## ABSTRACT

Internet of Things (IoT) devices create new ways through which personal data is collected and processed by service providers. Frequently, end users have little awareness of, and even less control over, these devices' data collection. IoT Personalized Privacy Assistants (PPAs) can help overcome this issue by helping users discover and, when available, control the data collection practices of nearby IoT resources. We use semi-structured interviews with 17 participants to explore user perceptions of three increasingly more autonomous potential implementations of PPAs, identifying benefits and issues associated with each implementation. We find that participants weigh the desire for control against the fear of cognitive overload. We recommend solutions that address users' differing automation preferences and reduce notification overload. We discuss open issues related to opting out from public data collections, automated consent, the phenomenon of user resignation, and designing PPAs with at-risk communities in mind.

## Author Keywords

Internet of Things; Personalized Privacy Assistant; Interviews

## CCS Concepts

•Security and privacy → Usability in security and privacy;

## INTRODUCTION

As everyday objects become internet-connected, with computational abilities beyond their original conception, they form an "Internet of Things" (IoT). As IoT grows [2, 34], new data collection risks and security exploits arise. From digital cameras to heart monitoring devices, the lack of security in such devices makes them fertile grounds for malware proliferation and privacy invasions [14]. Furthermore, poor disclosure to consumers about device capabilities and data practices [44] exacerbates privacy and security concerns [12, 30]. Existing notice and consent mechanisms are not easily or effectively implemented in IoT. Many IoT devices lack traditional interfaces

or user controls, and given the proliferation of IoT devices, informing users and asking for their consent to share data for every device they encounter can overwhelm them.

Privacy assistants (PA) have been proposed as a solution to help users manage their privacy in the face of increasingly complex and diverse data collection and use practices [5, 7, 26, 38, 41]. PAs can selectively notify users about data practices they would want to know about without overwhelming them [7] and also help users manage an increasingly large number of privacy decisions [7, 26]. In an IoT context where users are often not even aware of the presence of IoT devices and services collecting their data, PAs can also help discover IoT devices and services in the user's vicinity.

Researchers have identified factors that impact users' privacy decisions and comfort levels with certain data sharing practices [31], and have successfully built user profiles for privacy assistants [3, 25–27, 37, 42]. These findings allow personalized privacy assistant (PPA) designers to build decision-making models based on user preferences as a way to reduce user burden. However, this reduction in user burden can also reduce users' perception of control, increasing user anxiety and decreasing acceptance [40]. Striking a balance between autonomy and control is not a new issue [36] and has been explored from a theoretical perspective for context-aware systems [18] and privacy decision-making systems [6].

Interviews with users of a mobile app privacy assistant suggest that users are generally comfortable with the recommended changes to their mobile app permission settings [26], but other possible configurations of privacy assistants, including configurations that attempt to automate user decisions, could lead to different reactions. To the best of our knowledge, research on users' perceptions of different configurations of autonomy and control of PPAs for IoT scenarios is limited.

We examine the broader design space for configuring IoT PPAs. Specifically, we focus on understanding how users would respond to different possible PPA implementations that leverage predictions to assist users with privacy decisions. Examples of different implementations include: having PPAs check some of their conclusions with users, varying the extent PPAs can autonomously act on their own conclusions, adjusting the level of granularity at which PPAs allow users to reject/refine their recommendations, and how much users can refine these implementations.



This work is licensed under a Creative Commons Attribution International 4.0 License.

CHI'20, April 25–30, 2020, Honolulu, HI, USA.  
© 2020 Copyright is held by the owner/author(s).  
ACM ISBN 978-1-4503-6708-0/20/04.  
<https://dx.doi.org/10.1145/3313831.3376389>

We use 17 in-depth interviews to explore people’s opinions on PPAs. We examine end users’ perspectives on three hypothetical implementations of an IoT PPA, identifying perceived advantages and disadvantages of each implementation. The implementations vary along two of the four stages of the model of human information processing proposed by Parasuraman, Sheridan, and Wickens (2000): information analysis and decision/action selection. Specifically, we present three increasingly more autonomous PPA implementations: *Notification PPA* provides end users with awareness and control over data requests; *Recommendation PPA* gives users recommendations on how to respond to individual data requests; and *Auto PPA* makes autonomous data sharing decisions for the user.

We find that participants were excited about an aid that could provide recommendations from external sources, helping them make more informed decisions and bridging gaps in their knowledge. On the other hand, participants were less comfortable with a PPA being “too helpful” during the decision and action selection stage. These concerns mimicked participants’ overall concerns about IoT devices becoming too smart and people losing autonomy. Participants who were uncomfortable with a PPA that is “too smart” either rejected the Auto PPA or only wanted it to repeat decisions they had previously made (i.e., be automated, but not autonomous). Nevertheless, some participants viewed an autonomous PPA positively, as they trusted companies to make the correct decisions for them and appreciated that it would reduce the burden of privacy management. Throughout the interviews, participants weighed conflicting views of accepting more automation and reducing cognitive overload.

Based on our findings, we recommend avenues to implement a PPA for IoT that addresses the varying desire for control that users have while reducing demands on users’ attention. In addition, we present open issues raised by our participants related to managing consent withdrawal from public data collections, automated consent, the phenomenon of user resignation, and designing PPAs with at-risk communities in mind.

## BACKGROUND AND RELATED WORK

Parasuraman, Sheridan and Wickens (2000) proposed ten levels of automation and a four-stage model of human information processing: information acquisition, information analysis, decision and action selection, and action implementation [36]. We use this framework to guide our explorations into users’ preferences for varying levels of automation in different PPA implementations. In this section, we provide an overview of previous research on privacy assistants, highlighting the varying levels of automation in these implementations. We then present related work that explores users’ perspectives on PPAs with different levels of automation.

### Levels of Automation in Privacy Assistants

Privacy assistants (PAs) have been explored in the context of ubiquitous computing and IoT for many years—from the 2002 pawS, which leveraged machine-readable privacy policies to enable or disable sensors and devices based on user preferences [23], to the 2018 description of a privacy assistant aware of nearby devices and capable of selectively deciding when

to notify its user [7]. These assistants span a wide range of user involvement in the decision making process. At one end of the spectrum are PAs that serve as notice delivery mechanisms only, requiring users to make all of the decisions [19]. On the other end of the spectrum are PAs that make decisions for the user in either an automated or autonomous way. One example of an automated PA is the Dynamic Location Disclosure Agent (DLDA) [10] that enforces a default profile autonomously whenever faced with a context in which a privacy preference had not been specified.

Intermediate solutions rely on a high level of user involvement but pre-process information for the user. Examples of these solutions include systems that inform users about sharing actions in response to data requests [20] and present users with a pre-processed ranked list of available services that match their defined sensitivity levels [16]. These solutions attempt to reduce users’ cognitive burden, but still prompt them to make a decision on each occasion. Another type of solution further reduces user involvement by using machine learning models to decide when to notify users of data collection and even “semi-automatically configure privacy settings on the users’ behalf” [7]. This more personalized PA would avoid overwhelming users in an IoT world, but could also reduce user awareness and control. While automation in PAs can be used both in the process of notification as well as decision-making, we focus only on varying the level of automation related to the decision-making process.

### Opinions on Personalized Privacy Assistants

Extensive previous work has investigated how to implement a PPA—for example, by exploring the best way to build user preference models [3, 26, 31, 42] or how to present suggestions to users [11]. However, there is not as much literature that focuses on users’ perspective on and desire to use such tools.

Liu et al. (2016) explored a related question by soliciting participants’ opinions on using a PPA for smartphone apps through Likert questions in the exit survey of a user study on building user privacy profiles for mobile app permissions [26]. They found that recommendations were helpful, especially if presented in bulk, that having a permission manager allowed them to monitor apps, and that the time and delivery of nudges are important. Our work expands on theirs in two significant ways: first, by using interviews, we are able to obtain more detailed and nuanced opinions on PPAs, understanding not only people’s views but also their motivations; second, we present results on three variations of a PPA for IoT with different balances between user control and system autonomy.

Closer to our study, Zibuschka et al. (2019) studied users’ perspectives on an assistant for IoT that incorporates varying levels of automation (access control decisions, transparency, and location of processing). The authors combined an interview study with survey results to identify potential issues and concerns with this type of system and quantify the willingness to pay for such a system [47]. The more detailed picture of users’ opinions and concerns obtained in our study allows us to build on prior work and provide recommendations on how to build a PPA that may address issues faced by end users.

In a different application area, Namara et al. (2018) presented results of an interview study where they investigated users' perceptions of varying levels of automation on a personal assistant for Facebook privacy settings [32]. The automation levels used in that study—highlight, suggestion, and automation—are similar to ours. The corroborative and complementary nature of our findings suggests that opinions on the balance between user control and system autonomy for privacy assistants may be consistent across information technology domains.

## METHODOLOGY

We conducted 17 in-person semi-structured interviews with participants from Pittsburgh, PA, USA. In this section we describe the recruitment protocol, the interview procedure, and our analysis. The recruitment material, screening survey, and interview procedure are available as supplemental materials. The study was approved by our Institutional Review Board and the Department of Defense Human Research Protection Office.

### Recruitment

We recruited participants in two batches (nine before the 2018 holiday season and eight after) from the Pittsburgh metro area. We used Reddit, Craigslist, and posters at local library branches, bus stops, and coffee shops to recruit participants (18 years or older) for an interview study about their opinions and preferences about the Internet of Things. We made no reference to privacy or security to avoid biasing participants.

Twenty-eight potential participants responded to the recruitment ads and were invited to answer a screening survey. The survey contained questions about demographics and experience with technology, as well as questions related to technological knowledge, opinions about a connected world, and device ownership that were adapted from Mozilla's "Connected World" survey [30]. Twenty-four of the invited participants completed the survey. We used the screening responses to select a diverse sample of participants for one-hour interviews and contacted them via email to schedule the interview. Each participant received a \$25 Amazon gift card. We recruited and interviewed participants in small batches until variation in responses decreased significantly. We ultimately interviewed 17 participants.

### Interview Procedure

We conducted semi-structured interviews in a meeting space at Carnegie Mellon University. The same researcher conducted all of the interviews. A second researcher acted as a notetaker and asked questions at the end of each stage. Interviews were recorded and transcribed using an online service.

The interview protocol had three parts: exploratory, anchoring, and PPA. The goal of the exploratory part was to learn participants' opinions and understanding of IoT, while the anchoring part was to normalize participants' baseline knowledge of how IoT works. By the end of the anchoring part, if privacy had not been mentioned, we asked participants about their thoughts on data privacy as a way to engage participants in thinking about potential privacy issues. In the PPA part, we introduced the notion of a PPA for IoT, presented as a future project. We

explained that the PPA could identify both active data requests, such as a device requesting biometric data from the user's health tracker, as well as passive data collection, such as a smart device with microphone capabilities that could collect ones' utterances while in the vicinity of the device. We discussed three implementations of an IoT PPA with participants: Notification, Recommendation, and Auto.

#### *Notification PPA*

Notification PPAs can identify which devices requesting data are nearby and notify their users of their presence and requests, allowing users to allow or deny each request. This version provides users with full control over information analysis and decision selection. We used the following text to introduce participants to this concept:

An idea we had, and that we want to get people's opinions on before we build, is: what if there was an app on your phone that could tell you about the different types of data collections that are happening from the Internet of Things devices that are in the room or building that you are in. What do you think of this idea?

We did not initially offer participants the option to exert control over the different data collection practices happening around them. Nevertheless, once participants expressed their initial opinions about this implementation, we offered this alternative to see if their opinions would change.

#### *Recommendation PPA*

Recommendation PPAs build on Notification PPAs but provide users with suggestions on which data sharing decision to make based on their preferences. In this version, users still have full control over the decision selection but the system has more autonomy related to information analysis:

Now imagine this app could automatically recommend decisions for you based on your preferences on when to allow, deny, or limit different data collections. Would you use this feature?

#### *Auto PPA*

Auto PPAs would make data sharing decisions for the user. This would reduce users' cognitive burden but would also effectively remove their control from the process:

What if instead of just recommending decisions this app could, on its own, make this decision for you based on your preferences. Would you use this feature?

For each variation, we asked about their opinions and interest in this technology, including the underlying reasons. We also asked them to compare and contrast the three variations. Finally, we asked how they would like to interact with a PPA and which functionalities they would like to have.

### Analyses

At the end of each interview, both researchers wrote down their main observations before discussing the interview with each other. This allowed us to be aware of different perspectives and biases within the research team, and to identify topics that required further probing. After the first eight participants, we added questions at the end of the anchoring part to probe

participants about specific issues if they were not mentioned. All materials generated, including transcript and our structured notes were digitized. The research team reviewed the transcripts for correctness before coding them.

We coded the exploratory and anchoring sections separately from the PPA section. For all sections, two researchers went through the transcripts identifying passages that related to each of the questions from the structured notetaking form.

For the first two sections, one researcher conducted a first pass of open coding (or initial coding [39]) and produced a codebook with codes related to participants' opinions of IoT (positive, negative or neutral), positive and negative aspects of IoT, and their understanding of IoT. The researchers discussed and refined the codebook. Next, we performed a second round of coding where the main coder coded all aspects of the transcripts with another member of the research group as a second coder. We engaged in structural coding [39] for participants' valence of opinion of IoT and positive and negative aspects of IoT. We used holistic coding, evaluating all of the interview materials [39], for participants' understanding of IoT. Coders resolved conflicts and adjusted the codebook as needed.

For the PPA section, two researchers reviewed the passages related to participants' opinions about the different implementations of the PPA, whether opinions changed once more details were offered (i.e., when control was presented as an option for the Notification PPA), and interaction preferences. Each researcher summarized these passages into higher-level concepts, and they discussed those concepts together. After noticing a high level of consistency in the concepts identified, we opted to forego formalizing a codebook and iterative coding process, instead resolving the few conflicts that occurred.

### Limitations

Interviews are limited in terms of validity and reliability of the results [1]. We attempted to mitigate these issues by having one interviewer conduct all interviews to limit sources of inconsistencies, by having a dedicated notetaker who could serve as a separate check on the interviewer's understanding, by taking time to correct misunderstandings when the interviewer did not make themselves clear enough for the participant, and by avoiding leading questions. Furthermore, to avoid biasing participants, we did not mention privacy or security in the first half of the interview or in any communication prior to the interview, and we held the interviews in a building that is not associated with privacy or security research. We attempted to mitigate positivity bias [8] by stating that the PPA was something we were thinking about doing and wanted feedback before building it. When asked if we would build it ourselves, we stated that we would not. Even if our participants were somewhat positively biased, they still offered negative feedback. Nevertheless, the exploratory nature of the study and its small sample size are inherent limitations. We do not engage in comparisons between groups of users, as they would not be meaningful, and our results may not be generalizable.

### FINDINGS

In line with previous work, our participants had a good, but incomplete, understanding of IoT. They had concerns about

privacy, security, and safety, but also expressed concerns about losing autonomy and technology becoming "too smart." They expressed a lack of concern at times due to a focus on the societal benefits of IoT, expectation that companies are trustworthy, and nuanced feelings of resignation.

These viewpoints were later reflected in participants' opinions of PPA. We found that participants did not want pure awareness PPAs because they feared becoming overwhelmed and further resigned, wanting instead a system that gave them control over decisions. They viewed recommendations more favorably when seen as a way to bridge their knowledge gap as opposed to a paternalistic nudge, and they wanted to know and trust the sources of the recommendations. They had divided opinions on an Auto PPA due to conflicting desires to not be overwhelmed and to have control. The dislike of an Auto PPA was, at times, associated with a distrust that a system would be able to correctly predict their every decision, as well as a concern that IoT's societal benefits would be decreased if people could easily opt out.

### Participant Overview

Out of our 17 participants, eight self-identified as male, eight as female, and one as non-binary. Our participants had a mean age of ~39 years (min = 22, max = 68) and were fairly well educated (13 had some form of higher education). Most of them were employed (eight full-time, four part-time). Only three participants self-identified as students, and four reported working with or studying in areas related to technology or security. Only P10 and P11 were affiliated with our institution. Most of our participants self-identified as technologically savvy or experts, and were generally optimistic about a connected world. (A table with detailed demographics can be found in the supplemental materials.)

### Understanding of IoT

Almost all participants were able to describe the term "Internet of Things," although P1, P2, and P12 confused it with the Internet and P16 stated that they had no previous knowledge of the term. To evaluate participants' understanding of IoT, we asked them knowledge questions about five related concepts: types of devices, communication between devices, data collection, data storage, and data access. We scored participants' responses for the individual indicators and then generated the understanding categories by adding those values. We found that participants had an overall good knowledge of IoT. Please refer to the supplemental materials for a more detailed description of the categories and scores.

We identified knowledge gaps related to data storage, data collection, and device communication. Participants P1-2, P12, and P17 believed IoT devices store data only locally, while P5-9 and P14 described only a generic storage "cloud." Most participants anticipated either that devices would only store simple information (e.g., "on" or "off" states; P1, P3, P7, P10, and P12) or that devices would store at most usage logs (P8, and P13-17). Only P4 and P11-13 anticipated communication between devices. This seems consistent with findings that users underestimate what is being collected and potentially

analyzed by IoT devices [24] and that users struggle to understand privacy risks based on data inferences, especially risks posed by devices that do not record audio or video [46].

### Opinions About IoT

Nine participants viewed IoT as a positive direction for technology, as shown in Table 1. This aligns with previous work that found generally positive emotions associated with the term IoT [4]. Some of our participants stated that IoT was both positive and negative, but not neutral, as explained by P13: “It’s definitely not neutral. I think it’s good and bad and we don’t know yet which one it will be more of. . . . It’s going to be probably very good and very bad.” P4 perceived IoT’s present state as negative: “Bad under current trends.”

The most frequently mentioned benefit of IoT was its expected ability to make participants’ lives easier (e.g., convenience, saving time or money). This finding is similar to that reported in previous studies [12, 46, 48]. However, the second most frequently mentioned benefit was the opportunity to contribute to the greater good (P2-5 and P11-16), such as improved health-care and city planning, more innovation, resource conservation, improved emergency response, and avoiding accidents.

As in previous work, participants reported concerns about privacy, data security and misuse, and physical safety [4, 48]. With the exception of P11 and P14, participants mentioned privacy as a potential IoT issue unprompted. Furthermore, participants mentioned societal issues. Even though P10-12 and P14 saw potential benefits of artificial intelligence and autonomous devices, P1, P3, P5, P13-14, and P16 disliked the idea of machines becoming “too smart;” P11-12, P14, and P17 mentioned them taking jobs from humans; P2, P8-10, P13, and P17 were concerned about IoT data collection and inference leading to job or service discrimination; and, P1, P6, P13 worried about people becoming lazier. Some raised concerns about the impact on institutions and governments in the form of corruption (P4 and P13), manipulation (P5 and P11), and tailored echo chambers (P16). As P13 puts it:

It’s hard enough to maintain a sense of community and our democratic institutions are already really struggling . . . . To have all this sort of analysis going on behind the scenes to present you with something that, even if it is well intended, it’s still incredibly manipulative.

Participants who expressed non-negative opinions about potentially problematic scenarios mentioned common motivations, such as limited awareness of data collection and consequences [24, 46], and a focus on the potential benefits derived from IoT [45, 46]. Nevertheless, as P16 puts it, our participants saw societal benefits and benefits beyond the self as more worthy:

Like I would like to know what my data is being used for. . . . If it’s just for, you know, to recommend certain products to you, I find that a little annoying. Or if it’s something more useful, like, um, maybe it shows researchers how people stay connected. Maybe Facebook will be building more functions that help you stay connected to your friends instead of like pushing certain products.

Our participants showed more nuanced and complex motivations for nonchalance related to trust and resignation than previously found in the literature. P9, P13, and P15 believed they could protect themselves from privacy issues (e.g., by generally being ‘careful,’ by limiting automatic sync of devices, or by ensuring devices could be turned off), which has been previously found in the privacy and security literature [15, 22]. Some participants also trusted manufacturers [24, 43, 46]. However, this trust was articulated either in terms of an expectation that companies will do the right thing or in a conditional and hopeful way. P14 showed this expectation when he stated: “It’s not like a big concern because I believe that companies that manufacture the connected lock and connected cameras should have it very secure.” As P17 noted:

These companies, they are going to be big companies. They’re going to be ethical companies, you know. . . . They may invade my privacy, but I don’t see them doing much with it.

While some participants felt like they could protect themselves and trust companies, they did not always feel empowered by the options available to them. This came across as resignation, a phenomenon found in previous studies [24, 28]. However, our participants expressed feelings of resignation that went beyond mere powerlessness. We noted three different types of resignation: strong resignation, associated with expecting the worst of data collection and technology companies, marked as resigned in Table 1; normalization, an acceptance that “it is what it is;” and powerlessness to avoid a negative outcome.

P2’s comments on Google’s data collection practices is an example of strong resignation [9]: “Whether they do or don’t, or whether I try to stop it or not, that doesn’t really bother me, really. . . . *They’re gonna do it anyway. It doesn’t matter.*” The normalization of practices mirrored what has been found when people are forced to live with surveillance devices in their home [35]. P8 stated that “you can’t have a private conversation without the Internet listening. *So, it’s just more of that.*” Lastly, feelings of powerlessness stemmed from issues existing independent of IoT devices (such as home security) and due to an inability to use privacy and security solutions.

### Perceptions of Personalized Privacy Assistants (PPAs)

In this section we analyze participants’ perspectives on three implementations of PPAs for IoT. While we set out to explore specific implementations of PPAs, our participants also described their desired features and implementations. We present in Figure 1 the implementations with the details we envisioned in bold, namely notifications with and without control, recommendations based on previous behaviors, and an autonomous Auto PPA that leveraged users’ previous behaviors. We also include, in italics, additional implementations suggested by participants.

Participants’ opinions of PPA generally became less positive as automation increased, from Notification to Auto, but a few, such as P8 and P13, showed a non-linear relationship between their opinions and the level of automation. As we saw in the previous section, resignation was an important aspect, with P2

PID	Internet of Things		Data	PPA		
	Opinion	Understanding	Privacy Concern	Notification	Recommendation	Auto
P1	Positive	Low	Concerned	Positive +control	Negative	Negative
P2	Positive	Low	Resigned	Neutral	Negative	Negative
P3	Positive	Low	Resigned	Positive +control	Positive	Automated
P4	Negative	High	Concerned	Positive +control	Positive (education)	Autonomous
P5	Neutral	Average	Concerned	Negative	NA	Autonomous
P6	Both	Average	Resigned	Negative	NA	Autonomous
P7	Both	Low	Unconcerned	Neutral	Positive (education)	Automated
P8	Positive	Average	Neutral	Negative	Positive (education)	Negative
P9	Neutral	Average	Unconcerned	Positive +control	Positive (education)	Automated
P10	Positive	Average	Neutral	Positive +control	No opinion	Autonomous
P11	Both	Average	Neutral	Positive +control	Negative	Negative
P12	Positive	Average	Concerned	Positive +control	NA	NA
P13	Both	High	Concerned	Positive	Negative	Automated
P14	Positive	Average	Unconcerned	Positive +control	[Confused]	Automated
P15	Positive	Average	Unconcerned	Positive	Positive (education)	Autonomous
P16	Positive	Average	Unconcerned	Negative	Positive	Negative
P17	Both	Average	Concerned	Positive +control	Positive	Negative

Table 1. Participant characteristics identified during the interview.

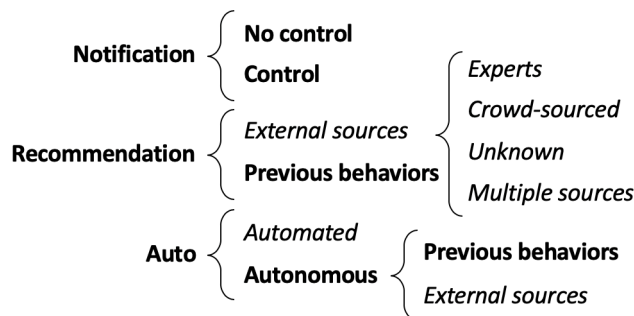


Figure 1. Diagram of the different implementations of PPA. In bold are the implementations we originally envisioned and in italic, the implementations suggested by participants during the interviews.

seemingly so deeply resigned that they did not see the purpose of privacy-protecting technologies.

#### Automation Level 1: Notifications

After completing the exploratory and anchoring sections of the interview, we asked participants about a Notification PPA that did not provide users with control over the data collection:

*“What if there was an app on your phone that could tell you about the different types of data collections that are happening from the Internet of Things devices that are in the room or building that you are in. What do you think of this idea?”*

Table 1 shows that most participants (n=11) had a positive reaction to an implementation of PPA that could provide users with awareness of data collection around them. However, this reaction was almost always accompanied by a desire, and at times expectation, that the system would also provide them with control over these data collections. P17 expected the following interaction when offered the above description:

The app would say... here are the smart devices within your range. ... It’s probably going to be where they connected to your phone. Here they are, here’s what they do... you can delete just one if you want. ... You can limit it if you want.

Only P13 and P15 did not expect or actively want control in order to accept this implementation of PPA. Nevertheless, when presented with the opportunity to exert control, they were happy to add this functionality. Another common expectation and desire was that the PPA should be easier to use and more efficient than existing privacy notices and policies. As P15 stated: “It sounds like a good idea to me... it sounds like a lot of boiler plate... like a terms of service thing almost. Although it’d be more concise and more specific, I guess.”

P5-6, P8, and P16 had negative opinions of this version of PPA. They were concerned that learning about all the information being collected about them would fuel their anxiety and feelings of despair. As P6 put it:

If I was always aware of [all the devices], I would be even more uncomfortable than I already am. It’s like, I don’t know where the line between knowing and just slowly spiraling into total paranoia, ‘cause I can see everything that’s being collected, all the time.

To prevent users from being overwhelmed by notifications, participants suggested only being notified about new and unexpected devices, or specific device types, being able to define “known locations” to not be notified by things in trusted spaces, receiving batched notifications, and being able to set the frequency of notifications. When we asked participants’ preferences on how the assistant should notify them, P1, P7, P9-11, and P15-17 preferred the push format, with P10 stating: “Push notifications would be much better, because I’m able to see all these notification instantly. I might not open the dashboard every time.” Nevertheless, P3, P4, and P13-14 liked having a web portal or dashboard where they could see a list of privacy

notifications on demand rather than relying on being able to access them in real time as they are pushed. P3 said “Either of those [dashboard or active push-notification] are a good idea. And then you have an option to not use up all your battery.” These participants also imagined that a dashboard would allow users to check their past decisions, and adjust and adapt as needed. P4 explains how this might work:

I guess having the option to do both. If we’re assuming that data collection is gonna be opt out rather than opt in, I would want a push identification . . . but then also be able to do kind of just an audit or a check in on my own, to just log in to this portal and see the full list and be able to switch things on or off at will.

Participants who had negative opinions about the Notification PPA that included control options also mentioned how this PPA could impact the societal benefits that could be derived from public IoT devices (e.g., cameras). In particular, they were concerned about allowing people to opt out of having their data collected from public-facing devices. One participant noted that someone could simply request that their data be deleted after committing a crime. P16 explained that allowing people to opt out could be problematic for smart cities:

I feel like that would really ruin a lot of the smart city sensors. If the point is to gather traffic information and you’re like, ‘no, you can’t count me.’ It will defeat the purpose of gathering data if you can opt out so easily.

P3 and P16 suggested that people should not be allowed to opt out of public-facing data collections. P3 noted that if you can tell cameras to not save your information “then people will just turn off cameras and rob everyone’s houses.”

These conflicting desires when managing data collection and opt-out were also noted when discussing private spaces. P8 wondered how conflicts would be managed if the Notification PPA gave users the ability to opt out of data collections—for example, one person might want a conversation recorded in the house, but the other one might not. They were particularly concerned about how this could impact women in situations of domestic abuse or contested divorces. They noted that an always-on listening device could gather evidence for those women. A PPA that allows people to opt out of data collection could limit this type of potential benefit. Finally, P6 suggested that the records of a user’s privacy choices could themselves create risk, e.g., if law enforcement became suspicious of what a person had chosen not to record.

Lastly, P2 and P7 did not see the usefulness of this tool, either for themselves or in general. At first P2 showed a neutral opinion towards this implementation of PPA and the data collection: “It doesn’t bother me.” However, when asked about PPA being able to opt-out of data collection, they showed how their resignation with the status quo impacted their motivation to protect themselves: “I think you can do that now. You can say it, but that doesn’t mean nothing. It’s gonna happen anyway.”

### Automation Level 2: Recommendations

Next, we asked about a PPA that could provide recommendations on whether to allow or deny a particular data collection:

*“Now imagine this app could automatically recommend decisions for you based on your preferences on when to allow, deny, or limit different data collections. Would you use this feature?”*

Opinions were, again, mostly positive—but there was a clear preference toward external recommendations (e.g., experts, manufacturers, friends) over recommendations based on past behavior. Overall, participants found this level of automation helpful and educational, and they appreciated that the tool could reduce their cognitive burden while augmenting their knowledge about what would be a good choice to make. P4, P7-9, and P15 explicitly stated their desire for this implementation of PPA to serve an educational purpose. As P4 stated:

When it comes to new technology, you can’t expect that people are gonna be perfectly literate or just be able to imagine all the ways in which information like that could be used. Now, for example, I guess you have a tool . . . and it will tell you, “This is the nature of this particular tracker. Here’s what it’s doing.”

We originally conceived and framed recommendations as based on users’ preferences. However, participants presented a range of opinions on possible sources. P4, P7, P9, P13, P15, and P17 suggested offering recommendations from authoritative sources with no vested interests in the data, but P17 also saw the benefits of recommendations from device manufacturers, since they know the technology the best. P7-9 suggested recommendations based on crowd-sourcing or user reviews.

Regardless of the source, participants wanted transparent disclosure of that source. P4, P7, P16, and P17 were aware that some sources of recommendations might have biases or even try to manipulate users. P16 also wanted to know the reasoning for the recommendation:

That’d be cool. But any sort of recommendation is biased. I would just wonder, how is that recommendation being made? Why is one thing more important than another?

P4 suggested allowing users to pick their preferred sources: “You can kind of get different information from different people, depending on what your own values are.” This could address the concern about the source of the recommendation. P7 suggested that having multiple sources could mitigate the sources’ biases and agendas: “It’s nice to have both perspectives. What do the security people who made this recommend? And then what the people who use it say about it?”

P1-2, P11, and P13 disliked the idea of a PPA providing recommendations and preferred to maintain control. P11 expressed concern about losing control over privacy decisions very clearly: “I should have the final say.” P1 framed recommendations as paternalistic or unnecessary. P1 said, “I think it’s simpler for it to just ask,” arguing that dismissing an unwanted recommendation would just add an extra step to her decision-making process.

*Automation Level 3: Automatic Decisions*

Lastly, we asked participants the following question:

*“What if instead of just recommending decisions this app could, on its own, make this decision for you based on your preferences. Would you use this feature?”*

About a third of participants did not want a PPA to make decisions for them. P1, P8, P11, P16, and P17 did not want to yield control over their decisions. P11 said, “I don’t like to be fully controlled by a device, you know?” Furthermore, P1, P8, and P17 were unsure whether the technology could accurately predict their decisions and were thus hesitant to allow it to do so. One reason for this, as P17 stated, is that we are not always consistent in our decisions: “I could change my mind. Nine times out of ten I’m going to go this way, but I’ve got a very good reason for that tenth time not to do that.”

Among the other two-thirds of our participants, we observed positive opinions towards a Auto PPA. These positive opinions reflected an appreciation of the convenience of outsourcing this type of decision-making to a computerized system. P5 rejected the Notification PPA due to a fear of becoming overwhelmed, but they said that the Auto PPA presented a valid tradeoff: “There I feel we’re obtaining a utility value to a human individual and I would consider owning such an appliance, as part of the digital world.”

P3-7 and P10 were happy to reduce their cognitive burden and potential anxiety from requests by delegating some of the responsibility of enforcing their preferences to the assistant. However, these participants were divided between an autonomous and an automated decision-making mechanism.

Initially, we envisioned this implementation of PPA behaving autonomously, leveraging users’ past decisions to build its prediction model and, when possible, making decisions for users without their involvement. On one hand, this implementation allows a PPA to still make decisions for the user when faced with new situations, by drawing inferences from similar past behaviors. This offers the lowest level of interruptions and cognitive load to the user, shy of a system that would completely remove the user. On the other hand, an autonomous PPA is prone to making incorrect inferences, especially if its decision model uses a feedback loop format where it uses (potentially incorrect) past decisions to inform future decisions.

About a third of our participants were comfortable having the Auto PPA function in an autonomous way (P4-6, P10, and P15). These participants trusted that decisions could be correctly made based on their trust and experience with other predictive technologies from existing companies. P5 related this to their experience with email—“[if] my email auto-complete capability is any indication, it can be pretty smart”—but added the caveat that “it would need to come from a company that I trust.” P15 highlighted another caveat, as they expected transparency and accountability:

That would be fine if there were a justification recorded somewhere where they would say on such date you made this decision. . . . I want to see the reasoning.

Citing similar concerns, some participants (P3, P7, P9, and P13-14) preferred an automated PPA—a PPA that would recognize when a user was faced with the same decision they had made in the past, and would implement that decision again—over an autonomous PPA. As P13 put it: “I think if they want to tell me ‘you’ve done this in the past, would you like to keep doing this?’ Yeah, that’s okay.” This version of the Auto PPA would offer fewer interruptions and lower cognitive load, while mitigating participants’ reluctance to give up control to a machine that might not make the correct decision. Nevertheless, this version would require more user input than an autonomous PPA and could still make wrong decisions for users if their data-sharing preferences varied over time.

P3-4, P6, P9 and P15 wanted an audit mechanism that they could use to see and correct previous decisions and to adapt preferences as they change over time. There was no consensus on whether changes to previous decisions should have retroactive implications or not, but P4 suggested: “I guess something useful to do would be like, have a system in which information is collected during a time period in which a preference has been set in an automated way, but you haven’t actually affirmatively audited it yourself.”

## DISCUSSION

By discussing different PPA variations, we noticed that participants frequently considered the trade-off between having more awareness and control versus feeling overwhelmed by data collection notifications. These considerations are in line with previous work exploring variables that may influence the desire for control and awareness while minimizing information overload [6]. Participants’ suggestions went beyond the levels of automation presented by Parasuraman et al. (2000), describing their desired implementation of automation for the PPA and the underlying approaches to achieve this automation.

Previous work found that users are more willing to purchase an assistant that has lower requirements of user attention [47]. In our work, we did not find that people’s opinions about PPA consistently increased or decreased with changes in user interaction or automation—some participants demonstrated a “curved” preference. This, plus the lack of a consensus on how to have more control without becoming overwhelmed, leads to a natural conclusion similarly found in previous work [32]: *a PPA needs to allow users to choose the level of automation they desire for the different functions being provided.* This allows users to balance their desire for control and their need to minimize cognitive burden. We present recommendations for promoting this flexibility in PPA design, focusing on the information analysis and decision selection functions [36].

### Designing the Information Analysis Function

We noticed that participants in our interviews were not always excited about automating the information analysis process. Participants were concerned about potential biases from the sources of recommendations or incorrect suggestions for inferred recommendations. Participants frequently adapted our prompt of a PPA that could infer patterns and propose recommendations on which actions to take, examining it instead as if the recommendations were from external sources.



One motivation for this change came from our participants' awareness that they had a good but incomplete knowledge of how IoT works. As such, they liked the option of a PPA that could educate them through external recommendations.

Participants had a wide range of opinions on what sources of recommendations were most preferred, wanted transparency about the source of the recommendations, and were aware of potential biases and agendas influencing these recommendations. Based on our findings, a good design for this functionality would allow users to **pick the types of recommendation sources** that they would prefer. Users wanted to see both expert opinions (authoritative) as well as real users' opinions (crowd-sourced). Social cues and expert recommendations can have a significant impact on users' decisions [11]. Furthermore, users have varied preferences and acceptance of different sources [33] and these sources have inherent biases. Because of this, a well-designed PPA **should allow users to choose preferred sources and should offer both crowd-sourced recommendations and recommendations from authoritative sources**. Furthermore, authoritative sources should include both manufacturers and independent organizations (e.g., Electronic Frontier Foundation).

The PPA must also ensure that users are not overwhelmed by notifications, not only because notifications can reduce individuals' ability to properly process the information being presented, but because notifications can make them anxious and resigned. One way to avoid this is to remove unnecessary notifications by **incorporating a "trusted location" feature**, such that users would not be notified about devices in those locations. This feature would not require significant effort from the user at setup, and it would avoid notifications for user-owned or known devices. For non-trusted locations, in lieu of potentially overwhelming individual notifications, the **PPA could list devices once and request a decision, revisiting that decision periodically** if new devices were added or the user's preferences changed.

Another way to avoid unnecessary notifications is to **specify data collection situations where users are always opposed to or always in favor of sharing**. Upon encountering these situations, a system could act automatically, saving notifications for situations where the user has not expressed a clear or consistent preference. Previous work on what factors people consider when thinking about IoT could provide the categories that users could select to always allow or always deny [12, 31].

Lastly, it is extremely important that the information be presented in a clear and informative way. Our participants sometimes felt resigned to or detached from privacy-concerning contexts because the risks did not seem tangible or did not offset potential larger societal benefits. To mitigate this, **recommendations should offer users tangible explanations of risks and benefits**. This could help users better understand and relate to the consequences of data collection, making them more likely to make appropriate decisions.

### Designing the Decision Selection Function

For the first level of automation (Notification), we observed a clear desire to have **a tool that not only provides infor-**

**mation about data collection, but that also collects and enforces users' preferences related to data collection**. While this form of control may not be ideal for many users—especially those that already have a tendency to normalize data collection or feel overwhelmed by current data collection—this basic level of control can help counteract feelings of powerlessness and can offer a starting point to further interactions with other modules of the PPA.

A next level of automation could implement users' predefined preferences in an automated way. For example, when prompted to make a decision, the user can choose to have the PPA **"remember this decision,"** informing the system that they no longer want to be asked about that specific data collection. This could be based on data types, devices, companies, etc. or a combination of these variables.

For decisions made without direct user input, it is vital to **provide an auditing mechanism** where users are able to verify and adjust decisions made on their behalf. This mechanism was considered essential when discussing an autonomous PPA, serving as a tool to avoid perpetuating incorrect decisions. This mechanism would also prove beneficial to users who choose a lower level of automation, as it would allow them to review and revise past decisions as their preferences evolve.

### OPEN ISSUES

The use of PPAs opens up a number of other questions fundamentally associated with the wider infrastructure and decision making that will influence the adoption of PPAs. In particular we describe questions related to public data collection, automated consent, resignation, and at-risk communities.

**Public Data Collection:** Allowing users to deny data collection might disrupt smart city functionalities and other technologies with broad societal benefit, if the proportion of non-contributing users grows past a certain threshold. Furthermore, safety and security devices could become useless depending on how this type of assistant is implemented. We propose the following open questions for further exploration: How do we balance the societal benefits of public data collection with individuals' desire for privacy? To what extent does anonymization, aggregation or differential privacy help mitigate people's reservations about some data collection practices and reduce the chance they opt out? What type of accountability system could be put in place to avoid ill-intended people bypassing safety and security devices?

**Automated Consent:** Previous work has demonstrated that automated consent models can predict users' data-sharing decisions with high accuracy and avoid prompting users for most decisions, drastically reducing user burden [42]. However, autonomous or automated consent is, at times, negatively perceived by users due to their desire to retain agency. Furthermore, if consent is given without the direct involvement of the consenting party, a legitimate question is whether it should be considered morally valid consent [21]. Participants did not trust their preferences to be correctly inferred since sometimes even they do not know what they would like to do until the very moment when they need to make that decision. One suggestion is to hold the data in escrow, so that it is not

immediately available to the requesting device, giving the user time to review it. Our questions then become: How would this system be implemented, and who would be responsible for the data in escrow? How would this system deal with time-sensitive requests? If a data escrow is not a viable solution, how can incorrect consent be corrected, and who becomes liable for the consequences?

**Dealing with Resignation:** Resignation is now noted as a common feeling associated with privacy [24, 28], but people displayed varying levels and types of resignation. Resignation stemming from powerlessness could be addressed through usable privacy-enhancing technologies, but a better understanding of causes of and solutions for resignation is crucial, as it can impact whether people attempt to engage in protective behaviors at all. While the identified level of privacy concern did not seem to play as large of a role in defining participants' opinions about PPA, P2 was so strongly resigned that they did not see the purpose of any sort of technological solution. Furthermore, those who had normalized current data practices showed even less interest in engaging with solutions that would require effort on their part. Legislation—such as the General Data Protection Regulation (GDPR) [13] and California Consumer Privacy Act [29]—could provide a baseline level of protection for heavily resigned individuals who are unlikely to engage with privacy-protective technologies. However, open questions remain: What causes the different types of observed resignation? What can be done to overcome these feelings of resignation so that people actively engage with privacy-protective solutions?

**At-risk Communities:** Smart home environments can raise issues of power dynamics, be that resident-visitor, parent-child, or between active and passive partner [17]. P12 and P15 mentioned potential benefits that a PPA could have on at-risk communities—for example, a person learning about surveillance from their abusive partner through discovery of IoT devices. However, participants also described potential risks. P6 expressed a concern that metadata about opt-out, which might include device location or time, could inadvertently disclose people's whereabouts and preferences. This could be particularly problematic for activists and other individuals who require anonymity. In these cases, open questions include: How do we ensure people's meta privacy in relation to their data sharing choices? How do we resolve preference conflicts when a preference impacts more than one party?

## CONCLUSION

We presented the results from 17 semi-structured interviews with a diverse sample of participants that allowed us to examine end users' perspectives on increasingly more autonomous possible configurations of a personalized privacy assistant (PPA) for IoT, identifying benefits and issues associated with each implementation.

Participants were generally positive toward the different implementations, however they also expressed concerns, which varied depending on the level of automation. Given the contrasting desires that seemed to drive participants—some wanted to have more control, while others wanted to avoid becoming overwhelmed by notifications—and the lack of consensus over

the best implementation of PPA, we recommend that PPAs be modular and offer configurable levels of automation, allowing users to select their preferred levels of control. We suggest features that will give users more control while still increasing the PPA's autonomy, such as letting users pick sources and have always/never scenarios pre-defined, as well as a mechanism for users to audit a PPA's decisions.

Finally, we discussed open issues that could impact the design and deployment of a PPA for the IoT. We formulate questions for future work that focus on managing consent withdrawal for public data collection, a way to mitigate the issues with automated consent, the phenomenon of resignation to the status quo of data collection, and designing these tools with at-risk communities in mind. Implementation of a prototype PPA that incorporates the suggestions made in this paper would allow for future studies with stronger ecological validity.

## Acknowledgments

We thank our reviewers and our shepherd, Dan Cosley, for their time and feedback. This research has been supported in part by DARPA and AFRL under agreement number FA8750-15-2-0277 and in part by NSF under grants CNS-1012763, SBE-1513957, and SBE-1514192. The US Government is authorized to re-produce and distribute reprints for Governmental purposes not withstanding any copyright notation thereon. Additional support has been provided by Facebook under an Emerging Scholar award and by Google. The views and conclusions contained herein are those of the authors and should not be interpreted as representing the official policies or endorsements, either expressed or implied, of DARPA, AFRL, NSF, Facebook, Google, or the US Government.

## REFERENCES

- [1] Hamza Alshenqeeti. 2014. Interviewing as a Data Collection Method: A Critical Review. *English Linguistics Research* 3, 1 (2014), 39–45. DOI: <http://dx.doi.org/10.5430/e1r.v3n1p39>
- [2] Consumer Technology Association. 2019. Spending on smart cities worldwide in 2015 and 2020 (in billion U.S. dollars). <http://www.statista.com/statistics/757638/spending-on-smart-cities-worldwide/>. (2019). Accessed: Feb 2019.
- [3] Paritosh Bahirat, Yangyang He, Abhilash Menon, and Bart Knijnenburg. 2018. A Data-Driven Approach to Developing IoT Privacy-Setting Interfaces. In *Proceedings of the 2018 Conference on Human Information Interaction & Retrieval - IUI '18*. ACM, 165–176. DOI: <http://dx.doi.org/10.1145/3172944.3172982>
- [4] Jiang Bian, Kenji Yoshigoe, Amanda Hicks, Jiawei Yuan, Zhe He, Mengjun Xie, Yi Guo, Mattia Proserpi, Ramzi Salloum, and François Modave. 2016. Mining twitter to assess the public perception of the Internet of Things. *PLoS ONE* 11, 7 (2016), 1–14. DOI: <http://dx.doi.org/10.1371/journal.pone.0158450>
- [5] Richard Chow. 2017. The Last Mile for IoT Privacy. *IEEE Security & Privacy* 15, 6 (nov 2017), 73–76. DOI: <http://dx.doi.org/10.1109/MSP.2017.4251118>

- [6] Jessica Colnago and Hélio Guardia. 2016. How to inform privacy agents on preferred level of user control?. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing Adjunct - UbiComp '16*. ACM, 1542–1547. DOI: <http://dx.doi.org/10.1145/2968219.2968546>
- [7] Anupam Das, Martin Degeling, Daniel Smullen, and Norman Sadeh. 2018. Personalized Privacy Assistants for the Internet of Things: Providing Users with Notice and Choice. *IEEE Pervasive Computing* 17, 3 (jul 2018), 35–46. DOI: <http://dx.doi.org/10.1109/MPRV.2018.03367733>
- [8] Nicola Dell, Vidya Vaidyanathan, Indrani Medhi, Edward Cutrell, and William Thies. 2012. "Yours is Better!": Participant Response Bias in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, 1321–1330. DOI: <http://dx.doi.org/10.1145/2207676.2208589>
- [9] Jillian D'Onfro. 2018. Google's Sundar Pichai was grilled on privacy, data collection, and China during congressional hearing. <http://www.cnn.com/2018/12/11/google-ceo-sundar-pichai-testifies-before-congress-on-bias-privacy.html>. (2018). Accessed: Feb 2019.
- [10] M. Elkhodr, S. Shahrestani, and H. Cheung. 2013. A contextual-adaptive Location Disclosure Agent for general devices in the Internet of Things. In *38th Annual IEEE Conference on Local Computer Networks - Workshops*. IEEE, 848–855. DOI: <http://dx.doi.org/10.1109/LCNW.2013.6758522>
- [11] Pardis Emami Naeini, Martin Degeling, Lujo Bauer, Richard Chow, Lorrie Faith Cranor, Mohammad Reza Haghighat, and Heather Patterson. 2018. The Influence of Friends and Experts on Privacy Decision Making in IoT Scenarios. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (nov 2018), 1–26. DOI: <http://dx.doi.org/10.1145/3274317>
- [12] Pardis Emami-Naeini, Henry Dixon, Yuvraj Agarwal, and Lorrie Faith Cranor. 2019. Exploring How Privacy and Security Factor into IoT Device Purchase Behavior. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM, Article 534, 12 pages. DOI: <http://dx.doi.org/10.1145/3290605.3300764>
- [13] EU GDPR Portal. 2017. Home Page of The European Union General Data Protection Regulation (GDPR). (2017). <http://www.eugdpr.org/eugdpr.org.html>.
- [14] IoT for All. 2017. The 5 Worst Examples of IoT Hacking and Vulnerabilities in Recorded History. <http://www.iotforall.com/5-worst-iot-hacking-vulnerabilities/>. (2017). Accessed: Feb 2019.
- [15] Alain Forget, Sarah Pearman, Jeremy Thomas, Alessandro Acquisti, Nicolas Christin, Lorrie Faith Cranor, Serge Egelman, Marian Harbach, and Rahul Telang. 2016. Do or do not, there is no try: user engagement may not improve security outcomes. In *Proceedings of the Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. USENIX Association, 97–111.
- [16] N. Gaud, A. Deen, and S. Silakari. 2012. Architecture for Discovery of Context-Aware Web Services Based on Privacy Preferences. In *2012 Fourth International Conference on Computational Intelligence and Communication Networks*. IEEE, 887–892. DOI: <http://dx.doi.org/10.1109/CICN.2012.52>
- [17] Christine Geeng and Franziska Roesner. 2019. Who's In Control?: Interactions In Multi-User Smart Homes. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM, Article 268, 13 pages. DOI: <http://dx.doi.org/10.1145/3290605.3300498>
- [18] Bob Hardian, Jadwiga Indulska, and Karen Henriksen. 2006. Balancing Autonomy and User Control in Context-Aware Systems - a Survey. In *Proceedings of the 4th Annual IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOMW '06)*. IEEE, 51–56. DOI: <http://dx.doi.org/10.1109/PERCOMW.2006.26>
- [19] Jason I. Hong and James A. Landay. 2004. An Architecture for Privacy-sensitive Ubiquitous Computing. In *Proceedings of the 2Nd International Conference on Mobile Systems, Applications, and Services (MobiSys '04)*. ACM, 177–189. DOI: <http://dx.doi.org/10.1145/990064.990087>
- [20] Hongxia Jin, G. Saldamli, R. Chow, and B. P. Knijnenburg. 2013. Recommendations-based location privacy control. In *2013 IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOM Workshops)*. IEEE, 401–404. DOI: <http://dx.doi.org/10.1109/PerComW.2013.6529526>
- [21] Meg Leta Jones, Ellen Kaufman, and Elizabeth Edenberg. 2018. AI and the Ethics of Automating Consent. *IEEE Security & Privacy* 16, 3 (may 2018), 64–72. DOI: <http://dx.doi.org/10.1109/MSP.2018.2701155>
- [22] Ruogu Kang, Laura Dabbish, Nathaniel Fruchter, and Sara Kiesler. 2015. "My Data Just Goes Everywhere." User Mental Models of the Internet and Implications for Privacy and Security. In *Proceedings of the Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*. USENIX Association, 39–52. DOI: <http://dx.doi.org/10.5555/3235866.3235870>
- [23] Marc Langheinrich. 2002. A Privacy Awareness System for Ubiquitous Computing Environments. In *Proceedings of the 4th International Conference on Ubiquitous Computing (UbiComp '02)*. Springer-Verlag, 237–245. DOI: <http://dx.doi.org/10.5555/647988.741491>
- [24] Josephine Lau, Benjamin Zimmerman, and Florian Schaub. 2018. Alexa, Are You Listening? *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (nov 2018), 1–31. DOI: <http://dx.doi.org/10.1145/3274371>

- [25] 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 *Proceedings of the Tenth USENIX Conference on Usable Privacy and Security (SOUPS '14)*. USENIX, 199–212. DOI : <http://dx.doi.org/10.5555/3235838.3235856>
- [26] Bin Liu, Mads Schaarup Andersen, Florian Schaub, Hazim Almuhammedi, Shikun Zhang, Norman Sadeh, Alessandro Acquisti, and Yuvraj Agarwal. 2016. Follow My Recommendations: A Personalized Privacy Assistant for Mobile App Permissions. In *Proceedings of the Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. USENIX Association, 28–41. DOI : <http://dx.doi.org/10.5555/3235895.3235899>
- [27] Bin Liu, Jialiu Lin, and Norman Sadeh. 2014. Reconciling Mobile App Privacy and Usability on Smartphones: Could User Privacy Profiles Help?. In *Proceedings of the 23rd International Conference on World Wide Web (WWW '14)*. ACM, 201–212. DOI : <http://dx.doi.org/10.1145/2566486.2568035>
- [28] Nathan Malkin, Julia Bernd, Maritza Johnson, and Serge Egelman. 2018. "What Can't Data Be Used For?" Privacy Expectations about Smart TVs in the U.S.. In *Proceedings of the European Workshop on Usable Security (EuroUSEC)*. DOI : <http://dx.doi.org/10.14722/eurosec.2018.23007>
- [29] Michael Morgan, Daniel Gottlieb, Matthew Cin, Jonathan Ende, Amy Pimentel, and Li Wang. 2018. California Enacts a Groundbreaking New Privacy Law. (June 2018). <http://www.mwe.com/en/thought-leadership/publications/2018/06/california-enacts-groundbreaking-new-privacy-law>.
- [30] Mozilla. 2018. 10 Fascinating Things We Learned When We Asked The World 'How Connected Are You?'. <http://blog.mozilla.org/blog/2017/11/01/10-fascinating-things-we-learned-when-we-asked-the-world-how-connected-are-you/>. (2018). Accessed: Feb 2019.
- [31] Pardis Emami Naeini, Sruti Bhagavatula, Hana Habib, Martin Degeling, Lujo Bauer, Lorrie Cranor, and Norman Sadeh. 2017. Privacy Expectations and Preferences in an IoT World. In *Proceedings of the Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)*. USENIX Association, 399–412. DOI : <http://dx.doi.org/10.5555/3235924.3235956>
- [32] Moses Namara, Henry Sloan, Priyanka Jaiswal, and Bart P. Knijnenburg. 2018. The Potential for User-Tailored Privacy on Facebook. In *2018 IEEE Symposium on Privacy-Aware Computing (PAC)*. IEEE, 31–42. DOI : <http://dx.doi.org/10.1109/PAC.2018.00010>
- [33] Bettina Nissen, Victoria Neumann, Mateusz Mikusz, Rory Gianni, Sarah Clinch, Chris Speed, and Nigel Davies. 2019. Should I Agree?: Delegating Consent Decisions Beyond the Individual. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM, Article 515, 13 pages. DOI : <http://dx.doi.org/10.1145/3290605.3300745>
- [34] Amy Nordrum. 2016. Popular Internet of Things Forecast of 50 Billion Devices by 2020 Is Outdated. <http://spectrum.ieee.org/tech-talk/telecom/internet/popular-internet-of-things-forecast-of-50-billion-devices-by-2020-is-outdated>. (2016). Accessed: Feb 2019.
- [35] Antti Oulasvirta, Aurora Pihlajamaa, Jukka Perkiö, Debarshi Ray, Taneli Vähäkangas, Tero Hasu, Niklas Vainio, and Petri Myllymäki. 2012. Long-term Effects of Ubiquitous Surveillance in the Home. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing (UbiComp '12)*. ACM, 41–50. DOI : <http://dx.doi.org/10.1145/2370216.2370224>
- [36] R. Parasuraman, T. B. Sheridan, and C. D. Wickens. 2000. A Model for Types and Levels of Human Interaction with Automation. *Trans. Sys. Man Cyber. Part A* 30, 3 (May 2000), 286–297. DOI : <http://dx.doi.org/10.1109/3468.844354>
- [37] Ramprasad Ravichandran, Michael Benisch, Patrick Gauge Kelley, and Norman Sadeh. 2009. Capturing Social Networking Privacy Preferences: Can Default Policies Help Alleviate Tradeoffs between Expressiveness and User Burden?. In *Proceedings of the 5th Symposium on Usable Privacy and Security (SOUPS '09)*. ACM, Article Article 47, 1 pages. DOI : <http://dx.doi.org/10.1145/1572532.1572587>
- [38] Norman Sadeh, Bin Liu, Anupam Das, Martin Degeling, and Florian Schaub. 2019. Personalized Privacy Assistant. (April 2019). Patent No. 20190108353, Filed Dec 29, 2017, Publication Apr 11, 2019.
- [39] Johnny Saldaña. 2009. *The coding manual for qualitative researchers*. SAGE. 223 pages.
- [40] J.G.M. van der Heijden. 2003. Ubiquitous computing, user control, and user performance: conceptual model and preliminary experimental design. In *Research Symposium on Emerging Electronic Markets*, U. Lechner (Ed.).
- [41] Junjue Wang, Brandon Amos, Anupam Das, Padmanabhan Pillai, Norman Sadeh, and Mahadev Satyanarayanan. 2017. A Scalable and Privacy-Aware IoT Service for Live Video Analytics. In *Proceedings of the 8th ACM on Multimedia Systems Conference (MMSys'17)*. ACM, 38–49. DOI : <http://dx.doi.org/10.1145/3083187.3083192>
- [42] P. Wijesekera, A. Baokar, L. Tsai, J. Reardon, S. Egelman, D. Wagner, and K. Beznosov. 2017. The Feasibility of Dynamically Granted Permissions: Aligning Mobile Privacy with User Preferences. In *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 1077–1093. DOI : <http://dx.doi.org/10.1109/SP.2017.51>

- [43] Peter Worthy, Ben Matthews, and Stephen Viller. 2016. Trust Me: Doubts and Concerns Living with the Internet of Things. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems - DIS '16*. ACM, 427–434. DOI : <http://dx.doi.org/10.1145/2901790.2901890>
- [44] Kaya Yurieff. 2019. Google says Nest Guard’s hidden microphone wasn’t meant to be a secret. <http://www.cnn.com/2019/02/20/tech/google-nest-microphone-secret/index.html>. (2019). Accessed: Feb 2019.
- [45] Eric Zeng, Shrirang Mare, and Franziska Roesner. 2017. End User Security & Privacy Concerns with Smart Homes. In *Proceedings of the Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)*. USENIX Association, 65–80. DOI : <http://dx.doi.org/10.5555/3235924.3235931>
- [46] Serena Zheng, Noah Apthorpe, Marshini Chetty, and Nick Feamster. 2018. User Perceptions of Smart Home IoT Privacy. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (nov 2018), 1–20. DOI : <http://dx.doi.org/10.1145/3274469>
- [47] Jan Zibuschka, Christian Zimmermann, Michael Nofer, and Oliver Hinz. 2019. Users’ Preferences Concerning Privacy Properties of Assistant Systems on the Internet of Things. In *Twenty-fifth Americas Conference on Information Systems*. AIS, 1–10.
- [48] 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*, Raimund Dachsel and Gerhard Weber (Eds.). Gesellschaft für Informatik e.V., 407–417.