Amazon vs. My Brother: How Users of Shared Smart Speakers Perceive and Cope with Privacy Risks

Yue Huang
University of British Columbia
Vancouver, Canada
huang13i@ece.ubc.ca

Borke Obada-Obieh
University of British Columbia
Vancouver, Canada
borke@ece.ubc.ca

Konstantin Beznosov
University of British Columbia
Vancouver, Canada
beznosov@ece.ubc.ca

ABSTRACT

With the rapid adoption of smart speakers in people's homes, there is a corresponding increase in users' privacy and security concerns. In contrast to previous studies of users' concerns about smart speakers' divulging private information to their manufacturers, our study focused on investigating users' concerns with regard to housemates and external entities. We conducted semi-structured interviews with 26 participants living in 21 households. Our results suggest that users often have an inadequate understanding of what data their smart speakers makes available to all users and what is kept private. Although participants expressed different privacy concerns about their housemates and external entities, they adopted similar, yet suboptimal, risk management strategies. We provide recommendations for future speaker design to support more optimal coping with the perceived risks.

Author Keywords

Shared smart speaker; Security and Privacy concerns; Mitigation strategies

CCS Concepts

•Security and privacy → Usability in security and privacy;
•Human-centered computing → User studies;

INTRODUCTION

A smart speaker is the most popular IoT device adopted at homes [49]. With it, users can carry out a number of activities such as requesting Uber rides, making purchases, and controlling other smart devices. Moreover, as Amazon has introduced more than 50,000 Alexa skills [4] into its smart speakers, additional features can be accessed, such as listening to TED Talks, playing games, and keeping up with sports news [31]. According to Voicebot's 2019 U.S. Smart Speaker Consumer Adoption Report [33], 26% of U.S. adults reported owning a smart speaker. Moreover, Loup Ventures predicts that 75% of U.S. households will have smart speakers by 2025 [32].

Although smart speakers have many desirable features, these devices have numerous privacy and security issues. Many studies and news reports have revealed vulnerabilities of smart speakers, such as weak voice authentication [38] and 24/7 listening and recording [28]. In addition, recently added features of Alexa (calling and messaging) have raised a number of security and privacy concerns among users. For instance, in May 2018, a Portland family stated in the news that their privacy was invaded when a private conversation in their home was recorded by Amazon's Alexa and sent to the phone of a person in the family's contact list [27]. Furthermore, it has been demonstrated that devices with voice interfaces can be controlled by voice commands that are unintelligible to human listeners [9]. Similarly, the high-frequency *dolphin attack* can trick speech recognition systems into receiving and responding to voice commands that cannot be perceived by humans at all [55]. Studies report that users' privacy and security concerns regarding the use of smart speakers include device hacking; recording of private conversations; 24/7 listening activities [41]; the collection, sharing, and storage of private data [1, 40]; and the "creepy" nature of the devices [41].

At the same time, security and privacy concerns due to sharing a smart speaker among *multiple users* in a single household have received little attention. Zeng et al. [54], Lau et al. [35], and Geeng et al. [19] discovered *power imbalances* between primary and secondary users. Similarly, Yao et al. [53] found that users considered social relationships and the power dynamics among multiple users when designing privacy mechanisms for smart homes. A survey conducted by Malkin et al. [40] discovered that of those who reviewed previous smart speaker actions, 56.8% admitted to encountering others' recordings while doing so. Many smart speaker companies have now enabled multiple users to set up their personalized accounts and voices on one smart speaker (e.g., [24]).

Since more than one-quarter of U.S. households now own a smart speaker [33], there is a need to explore users' privacy and security concerns in this scenario and provide support to improve their experiences. Many studies simply focused on investigating users' concerns with respect to external entities,¹ and few shed light on users' concerns about others living in the same household. We aim to bridge this knowledge gap

¹In this paper, external entities refer to parties that are not within the household environment, such as the speaker manufacturers, the government, and hackers.

by studying and comparing users' mental models, privacy concerns, and coping strategies as directed toward entities both inside and outside the household. An exploration of users' experiences with shared smart speakers could help researchers discover the underlying issues with the current design of smart speakers, thereby informing the future design of these devices.

We conducted semi-structured interviews with 26 participants from 21 households. Our results suggest that participants had limited understandings of how the smart speakers were shared among users. They also identified different security and privacy concerns regarding housemates and external entities. The main concerns identified by the participants about their households were voice match false positives, inappropriate access to personal information, and the misuse of the device by unintended users (e.g., visitors). Contrary to previous findings, the major concern specific to external entities was the collection of data by the smart speaker manufacturers for questionable purposes. Avoidance and acceptance were the major strategies adopted by the participants to address their concerns regarding housemates and external entities. Based on our findings, we recommend that future speakers provide a personalized sharing function, improved recognition of users' voices, and more effective delivery of technical support.

In summary, our contributions are as follows. We investigate users' understanding of how smart speakers are shared among multiple users within a household. We identify users' security and privacy concerns when making use of shared smart speakers in a multi-user scenario and their concerns toward external entities. We discover the mitigation strategies that users employ to address their concerns.

RELATED WORK

In this section, we first summarize the research that has explored users' privacy and security concerns about using IoT devices and smart speakers. Then, we review the reports of these risks that exist in the IoT domain and the vulnerabilities associated with smart speakers.

Privacy and security concerns about IoT. Prior studies focused on understanding users' perceptions [51], privacy and security concerns [54, 56, 14], and mitigation strategies [2] in the context of home IoT devices. Tabassum et al. [51] report that users are highly uncertain about the data practice of their smart home device manufacturers. Zeng et al. [54] suggest that users' threat models have gaps that arise from a limited technical understanding of smart home technologies. Therefore, it is not surprising that participants in a study by Naeini et al. [13] were significantly more concerned with their biometric data being collected and used by companies. These limited concerns might be in part explained by the findings of Zheng et al. [56] that owners of IoT devices value the convenience and connectedness over their privacy. When Aleisa et al. [2] conducted a literature review of existing IoT privacy-preserving solutions, they concluded that most solution providers assume that the end users would be willing to expend effort to preserve their privacy, which may be unrealistic.

Other scholars have studied users' privacy preferences when using IoT devices. Lee et al. [36, 37] discovered that while

participants were comfortable with device manufacturers' one-time monitoring of their activities, they were concerned about continuous monitoring. Regarding data access, participants were uncomfortable if their collected data were being monitored and shared with the government [37, 56]. He et al. studied participants' preferences for access control policies for different capabilities within a single IoT device. They discovered that the desired policies vary widely based on the relationship between the users and the context of access [26].

Privacy and security concerns about smart speakers. Lopatovska et al. [39] investigated the types of tasks that participants typically request from smart devices. They found that speakers were mostly used to obtain weather forecasts, play songs, and control other devices. Although Amazon claims that it only saves the audio of speech immediately following Echo's *wake word* [12], a 2014 patent application suggests that the device could also listen for a list of words that indicated statements of preference [11]. Abdi et al. [1] studied users' perceptions about 4 main use cases regarding smart speakers. They discovered that most users had limited conceptions of the smart speakers ecosystem and related data activities. As reported by Lau et al. [35], smart speaker users trust the speaker manufacturers to protect their privacy. In contrast, many non-users do not believe that the speaker company will collect only their voice commands and keep the data confidential. Furthermore, Chung et al. [10] investigated the possibility of predicting users' personality and daily lifestyle choices based on the types of requests made to Alexa. The authors discovered that they could easily determine users' sleep routine, dinner time, driving routes, and general interests based on the data that Alexa stored in the cloud. The above research was focused on investigating users' concerns in regard to entities outside a household, such as a speaker manufacturers and hackers. Since smart speakers are mostly adopted in home environments [49], exploring users' concerns about other people in the same household could inform the design of future generations of smart speakers.

Risks associated with the IoT. Security researchers have been actively investigating the vulnerabilities and security challenges in the IoT domain. Geneiatakis et al. [20] discovered vulnerabilities that enable impersonation, denial of service attacks, and eavesdropping on commands, while Bugeja et al. [8] identified possible vulnerabilities in services and communications of IoT devices. In addition, Bugeja et al. [8] reported four significant privacy and security challenges in the smart home domain, including identity management, risk assessment methods, information flow control approaches, and security management methods.

Risks associated with smart speakers. In relation to the privacy and security vulnerabilities of smart speakers, Jackson et al. [28] studied Amazon Echo features and discovered several vulnerabilities. The major one was the inability of the Echo to distinguish between voices or to detect a physical presence, making it easy for attackers to say the *wake word* and issue their own requests. Lei et al. [38] also explored the vulnerabilities of the Alexa-enabled Echo. They identified three vulnerabilities related to insecure access to the device:

weak single-factor authentication, no physical presence-based access control, and insecure access control on the device cloud. Similar to Carlini et al. [9] and Zhang et al. [55], Roy et al. designed an ultrasound signal that was demonstrated to remain inaudible to humans but was recordable by unmodified off-the-shelf microphones [47]. Furthermore, Roy et al. built on this principle to proposed a technique for long-range voice command attacks on smart speakers [48]. Gao et al. proposed a framework that could continuously jam voice-based assistants to prevent them from recording users' speech unless the user issues a voice command [18]. Additionally, Alhadlaq et al. [3] discovered that most Alexa skills did not have privacy policies, and, that of those that did, only 3.5% possessed a valid policy. As a proof of concept, researchers at Checkmarx developed an Alexa skill that could eavesdrop on users [45].

Our research contrasts with prior work in three major ways. First, as opposed to studying *all* IoT devices, we focus on smart speakers. Second, we study users' security and privacy concerns in a multi-user scenario, in which smart speakers are shared among housemates. We also investigate users' concerns about external entities and their corresponding coping strategies. Finally, we compare the coping strategies of users toward their housemates and external parties.

METHODOLOGY

Recruitment and Screening

Participants were recruited by advertising on Facebook, Twitter, Instagram, and UBC's paid participants studies list. Potential participants completed a survey to verify their eligibility for the study. The survey included questions to identify whether they were sharing their smart speaker(s), and the relationship between them and the other users. The screening survey also ensured that the sample covered a wide range of relationships within a household, such as roommates, parent-child, siblings, and romantic relationships. Based on the participants' preference, we conducted interviews either in person or via a video link. The participants were compensated with either \$20 in cash or (in the case of video interviews) Amazon gift cards. This study was reviewed and approved by the UBC Behavioural Research Ethics Board (ID: H18-01943).

Interview Procedure

Semi-structured interviews were conducted by two researchers, with one leading the interview while the other took notes and audio-recorded sessions. The interview questions focused on four segments (see supplementary material for details):

General questions: Participants were asked to describe their smart speakers, who they share the speakers with, the activities carried out on the device, the factors that influenced their device adoption, and their general disposition toward the device.

Mental model of shared smart speakers: To gain insights into users' perception of smart speakers, we used a combination of two methods: a drawing exercise and a verbal explanation. Drawing has been found to be a complementary method to the verbal reports used to understand users' mental models [29] and has been applied in many usable security studies [54, 46, 30, 16]. To avoid putting ideas into the participants' minds, we first asked each participant to explain how a

smart speaker works by drawing a diagram. After each drawing was completed, we asked participants to verbally describe their thought process with reference to their drawings. If there was any confusion, follow-up questions were asked to clarify the participants' drawings and verbal descriptions. If the drawings did not consider multi-user scenarios, the participants were asked to draw a separate diagram to explain how they perceived smart speaker functions in a multi-user scenario.

Security and privacy concerns about housemates and visitors: We asked questions that explored participants' security and privacy concerns with respect to the people they shared the device with. For instance, we asked the participants which features they believed could be accessed by all users and which could be used only by them. If the participant voiced a concern, we further investigated the coping strategies they used to address that concern.

Security and privacy concerns about external entities: We asked the participants if they had any hesitations before purchasing the device, if they had heard of any smart speaker-related negative news, and how they perceived this news. Furthermore, we asked the participants about their concerns and perceptions of the attacks or vulnerabilities that they mentioned. We also asked participants who they believed the adversary was in the context of these attacks or vulnerabilities. When asking questions about their concerns, we did not mention any specific threats to avoid biasing their responses. When a participant expressed a concern, we asked if they had employed any strategy to address that concern.

Data Analysis

The interview data were analyzed using Grounded Theory [43, 21]. After each interview, two researchers transcribed and analyzed the collected data. Based on the findings, the interview guide was modified before the next interview took place. The steps taken to analyze the data included open, axial, selective, and theoretical coding. During open coding, we identified 238 codes. In axial coding, we related the codes to one another, which resulted in the creation of 7 categories. All 3 authors worked together to select the main category and reorganize the related codes. We then developed the theoretical model to explain the core category, namely, the participants' mental models of smart speakers, their perceived concerns with respect to the shared smart speakers and external entities, and their corresponding coping strategies. Theoretical saturation was reached after interviewing 23 participants. We conducted 3 more interviews and did not obtain any new codes.

RESULTS

We conducted semi-structured interviews with 26 participants from 21 households. Their ages ranged 19-55 years (mean 31, median 30). Their demographics and other details are summarized in Table 1. Seventeen participants shared their device(s) with their family members, such as partners, siblings, children, and grandchildren. Others shared devices with friends. Twenty-three participants shared only one smart speaker in their households. Twelve participants were interviewed in person, while the other 14 were interviewed using video calls. Seventeen participants were identified as primary users based

P#	Device type	Age	Gender	Occupation	User type	Share the device with	Education level
1	Google Home	36	F	Project manager	Primary user	Father	Master
2	Google Home	31	M	Software engineer	Primary user	Wife	Bachelor
3	Google Home	30	F	Accountant	Primary user	Husband	Bachelor
4	Google Home Mini	25	M	HR assistant	Secondary user	Roommate	Bachelor
5	Google Home Mini	25	M	Unemployed	Primary user	Roommate	Bachelor
6	Amazon Echo	25	M	Master's student	Secondary user	Roommate	Bachelor
7	Amazon Echo Dot	29	M	PhD student	Primary user	Roommate	Master
8	Google Home Mini	20	M	Undergraduate student	Primary user	Brother	High school
9	Amazon Echo Dot	22	F	Undergraduate student	Secondary user	Sister	High school
10	Google Home	19	M	Undergraduate student	Secondary user	Brother	High school
11	Amazon Echo	20	M	Worker at an industry supply company	Secondary user	Sister	High school
12	Apple HomePod	38	M	IT technician	Primary user	Daughter, stepdaughter, granddaughter, and wife	Bachelor
13	Google Home Mini	32	F	Social worker	Secondary user	Husband and daughter	Bachelor
14	Google Home Mini	33	M	Film and television producer	Secondary user	Wife	Bachelor
15	Google Home	38	M	Nurse educator	Primary user	Son	Bachelor
16	Google Home Mini	38	F	Librarian	Primary user	Husband	Master
17	Amazon Echo (2nd Generation)	24	F	Research coordinator	Primary user	Roommate	Bachelor
18	Amazon Echo (2nd Generation)	25	F	Software developer	Secondary user	Roommate	Bachelor
19	Google Home Mini	38	M	Marketing	Primary user	Wife and daughter	Bachelor
20	Google Home Mini	27	F	Saleswoman at Best Buy	Primary user	Roommate	Bachelor
21	Google Home	25	F	Respirator therapist	Primary user	Mother and brother	Bachelor
22	Google Home	Prefer not to say	M	Consultant	Primary user	Roommate	Bachelor
23	Google Home Mini	55	M	Retired	Primary user	Wife and children	Bachelor
24	Google Home Mini Google Home	41	F	Compliance officer at financial services	Secondary user	Children	Bachelor
25	Google Home Mini Amazon Echo Amazon Echo Dot	40	F	Marketing consultant	Primary user	Husband and daughter	PhD
26	Google Home Mini Google Home	41	F	Technical support at telecom	Primary user	Mother and sister	Bachelor

Table 1. Summary of the participants' demographics

on a set of factors such as who took the main responsibility for setting up the device, who had the authority to change the settings, and who usually explored the new features and informed others about them. The interviews lasted an average of 40 minutes. The inter-coder agreement was calculated as 89%, a high level of agreement between the two coders [22].

Mental Models of Smart Speakers

Similarly to the findings of Tabassum et al. [51] and Kulesza et al. [34], we discovered that our participants' mental models of the smart speakers could be categorized into two different levels of sophistication. Advanced mental models (similar to the *structural* models in [34]) indicate that users have an in-depth understanding of how the device works, whereas limited mental models (similar to the *functional* models in [34]) imply that the end users only know the functions of the device. Our analysis was based on codes that indicated whether a participant demonstrated an understanding of smart speaker architecture, including data flow, processing, and storage.

Participants with advanced models ($n = 9$) had a highly technical understanding of their smart speakers and were able to represent the network topology, including the cloud servers, wireless protocols, and sometimes the routers. For example, P2 drew an accurate and detailed representation of the smart speaker architecture and explained how the data were stored and processed (see Figure 1a). All the participants with advanced models discussed how the data flow between the device and servers in the cloud during interactions.

Participants with limited mental models ($n = 17$) had some sense of how the smart speaker worked but were not well

aware of the technical elements. They were more focused on describing the services that the device could provide and did not have technical knowledge of how data flowed, were processed, or stored (see Figure 1b).

Mental Models of Shared Smart Speakers

Based on participants' drawings and verbal explanations of their **shared** smart speakers, we also categorized the sophistication of participants' mental models into *advanced* and *limited*. Participants with advanced models ($n = 6$) had a reasonable understanding of what features or data were shared and which were kept private, regardless of whether the features were currently in use. For example, P2 and his wife set up the voice match feature of their shared Google Home. P2 correctly explained that his contact list, calendar, and reminders were kept private and could only be accessed using his voice (Figure 2a). Another example is P16, who, even though she did not set up the voice match feature, was able to verbally explain that her contact list was available to all users.

Participants with limited mental models ($n = 20$) had an incomplete understanding of which features and data were shared and which were kept private, even after direct prompting. For example, when P22 explained how the smart speaker was shared between him and his family members, he only drew and described the features that were mostly used by each user. When we asked him about data sharing among users, P22 admitted to having no clue about this aspect (Figure 2b).

As discussed in the next section, the degree of sophistication of a participant's mental model of a shared smart speaker appears to correlate with their level of concern about their housemates.

Figure 1. Participant drawings showing examples of (a) advanced (from P2) and (b) limited (from P6) mental models of smart speakers.

Figure 2. Participant drawings showing examples of (a) advanced (from P2) and (b) limited (from P22) mental models of shared smart speakers.

Concerns about Housemates and Visitors

Concerns Based on Conjectures

Participants' inadequate mental models may lead to inaccurate threat models. With Google speakers, if no user has linked their voice to the device, everyone in the household can access the personal contact list of the first user to link the device to their Google account [25]. Similarly, by asking for a particular contact's name from shared Alexa-enabled devices, anyone can access the contacts of users in the Amazon Household. Alexa will, in turn, search for the requested contact on all available contact lists and make a call via that contact's profile [6]. Users can prevent the sharing of contacts by setting up their individual voice profiles. P19, who did not set up voice profile, expressed concerns about his 8-year-old daughter using the *redial* function of the device: "If I called [my boss], or if I called a client, I do not want my daughter [redialing] the last call to someone she is not supposed to talk to." However, the *redial* function is only available when the voice profile is set up. This misunderstanding confirms that P19 had an inadequate mental model of the device sharing, leading to an inaccurate threat model.

P12 was in a similar situation. His Apple HomePod did not offer the option of creating an individual profile to keep the information of each user private at the time [42]. P12 stated

his concern: "[My housemates] could find out what I have on my contact list ... it could be an issue as well." He further explained that "... they could have a history of who I recently called." However, P12 could not explain how his housemates could gain access to his calling history, which makes his concern more of a conjecture. Since HomePod did not provide a way for users to check the call history, P12's concerns suggest that his incomplete mental model of how the device is shared led him to an incorrect threat model.

Concerns Due to the Immaturity of the Technology

Participants who had an advanced model of how speakers were shared had a reasonable understanding of what data were shared among users and what data were kept private. As a result, they tended to express more concerns related to immature technology, such as voice match false positives.

Unauthorized voice purchases: Several participants indicated that, after setting up the voice profile feature, their smart speaker(s) occasionally failed to distinguish their voices from those of other users. This led to some ($n = 3$) participants being concerned about other users impersonating them and making unauthorized voice purchases. P8, who had a good understanding of how Google Home Mini was shared between him and his brother in terms of feature sharing, illustrated

his concern as follows: *“I do not feel the [purchasing feature is] secure enough ... because my younger brother has a similar voice. Sometimes when we talk ... [to] the speaker, it recognizes him as me. So [the speaker] may mess things up.”*

Unauthorized access to calendars and reminders: There were concerns about unauthorized access to participants’ personal information by housemates. Some ($n = 3$) participants with advanced models of shared speakers were worried that others could impersonate them to access their private information. P22 estimated that his Google Home did not recognize his voice 25% of the time, which caused him to worry that his roommate could access his calendar. P10 had similar concerns: *“My voice and my brother’s voice are similar ... I do not trust [the voice match of Google Home] yet. And someone who has a deep voice or someone [who] has a voice changer ... could use my voice to get access [to my reminders].”*

General Concerns about Other Users

The following concerns were raised regardless of the sophistication of participants’ mental models of smart speaker sharing.

Overheard call conversations: There were concerns about housemates overhearing phone conversations conducted over the speaker. These concerns were largely due to the fact that, when making phone calls using the speaker, the responses given by the speaker can be overheard by people nearby. This concern indicates that the calling feature is not always suitable for use while others are present, suggesting that the design of this feature does not take users’ privacy into account. P23 justified his concern: *“I do not want other people in the household to hear me talking about work, and my wife does not want everyone else to hear her talking to her friends.”*

Misuse by unintended users: Participants were concerned with the possibility of both, (1) young children living with the participants and (2) visitors, making use of smart speakers. To illustrate, P19 commented: *“[As for] the contact list, ... It is my [8-year-old] daughter that I am worried about.”* Our participants also expressed concerns about visitors, including their friends and less familiar people, such as party attendees. P16 expressed her concerns about party guests: *“If I have a party with a whole bunch of acquaintances or strangers, I unplug it and put it away ... because if I leave [the speaker] out for strangers, [they] might steal our WiFi, and then weird things [will] happen.”*

Security and Privacy Concerns about External Entities Company Data Collection and Usage

In terms of the concerns regarding external entities, many participants ($n = 13$) admitted to being mainly influenced by social media. Similarly to Tabassum et al. [51], we could not find any differences between the participants with mental models of different sophistication levels.

A popular concern identified by the participants was the collection of data by speaker manufacturers. Consistent with previous findings [56, 17], our participants were also concerned about the uncertainty and scope of the data being collected. Unique to our study is the finding that participants were concerned with the **usage of the collected data** by the speaker company. These concerns indicate participants’ lack

of trust in the speaker manufacturers, although they did not have evidence of suspicious behaviors by the companies.

Data used to determine life patterns: P24 was concerned about the device manufacturers using the collected data to determine her family’s life patterns. She explained that based on the family’s use of the device, the company could potentially learn about her family’s whereabouts within the home, and the times they were at home, which was verified by previous study [10]: *“Obviously, we have to be home when it is being used, so Google is able to establish the pattern of when we are home ... If I am asking for recipes, I must be in the kitchen.”* More specifically, a few participants ($n = 3$) were convinced that the speaker company could track their online behaviors. Similar to the findings in [1], our participants also believed that the company could review their search history and audio files with the aim of providing more targeted advertisements to them. To illustrate, P14 said: *“[The collected data] makes it easier for [the company] to target me with advertisements and track my behavior. I do not want to be more tracked than I already am.”*

Selling data: Participants believed that the speaker manufacturers could secretly sell their data to third parties. Several participants ($n = 6$) referred to Facebook’s data privacy scandal [7] to explain their concerns about the speaker company doing something similar. P16 described his suspicion of Google: *“Facebook was selling personal data to companies without users’ knowledge [and] was promising privacy, then giving full access to companies, even if the company did not know they had it. [Google] is a different company, but it is still in the same marketplace, that is the part I am worried about [Google selling my data].”* Furthermore, participants argued that the company sold their data for its own benefits rather than to benefit users: *“I think other companies could pay Amazon to access information from Amazon’s cloud”* (P9). P22 accused smart speaker manufacturers of being dishonest about their policy and secretly selling users’ data: *“I think [the smart speaker companies are] lying, because they say they need to use our information to make [the smart speaker] smarter, but I do not think [that is what they use the information for] ... they sell it.”* Furthermore, participants believed that selling their data was a serious violation of their privacy. P15 said: *“I do not like [my data] being sold, because you would wonder about how much personal information they have.”*

Conversation recording and sharing: Contrary to the findings by Tabassum et al. [51], a few of our participants ($n = 4$) were indeed concerned about their private conversations being recorded and shared with other entities. For example: *“I don’t know if my conversation gets transmitted or [not] ... whom the information is shared with would be the biggest concern ... [I am] not feeling [that I] have got privacy in my own home.”* (P13). Referring to the recent news report [27], P23 explained: *“It hit the news over the past 6 months, about smart speakers accidentally recording the conversation and playing it back for other people. I do not think that was very good, and I have concerns about how it could possibly happen.”*

Unclear data collection: Participants were concerned about the company’s right to collect their data. A few participants

($n = 3$) suspected (correctly [11]) that the company collected their data without their permission or collected more than it needed to. Others commented that they got tricked into giving permission to the company by blindly agreeing to the companies' terms of services. P10 explained his doubt: *"I feel like I did not really give [the smart speaker company] explicit permission. But when you press accept [to the terms of services, the speaker manufacturers] kind of trick you. Like, I consent to it in a tricky way ... because who reads the user-term guidelines?"*

Unclear trade-off between functionality and privacy: Similarly to the findings of Zeng et al. [54] and Zheng et al. [56], our participants explained that they had to accept the privacy and security risk in exchange for the functionality of the smart speaker. What we found *new* is that a couple of participants expressed concerns about the **uncertainty** of such a trade-off. For instance, P23 expressed his astonishment when he found out that his previous online shopping history was available on the Google Home App [23], though he had not used the speaker for purchases: *"I clicked on the purchase as part of [Google Home App] and [Google Home App] told me all the things I bought, that I did not buy through [the Google Home Mini]. I do not like how [Google Home Mini] went through all my emails and found all my [previous] purchases for me."*

Coping Strategies

Strategies Regarding Housemates and Visitors

Although participants had different levels of understanding of how their speakers were shared, they adopted similar strategies for coping with their concerns about housemates and visitors:

Avoidance: Some ($n = 11$) participants refrained from using the purchasing feature of the device to avoid possible voice match false positives and the misuse of the device by their young children. P8 described his coping strategy as follows: *"Because my younger brother has a similar voice [to mine] ... I think that is one of the reasons I do not want to link my credit card or any payment kind of stuff. Because [he] may access it and buy stuff."* Others did not use the calling feature because of their concerns about their housemates overhearing their conversations and accessing to their contact list. In addition, several participants reported putting the device away as a strategy to mitigate their concerns about visitors.

Acceptance: Even though many participants were aware of the potential privacy and security issues, they were not explicitly concerned about them. Instead, participants expressed acceptance of the risks, voicing several reasons:

Trust toward housemates: One reason for acceptance was that participants trusted that housemates would not act maliciously (even if they could), because there were no benefits to be gained. P11 justified his acceptance of the risk: *"I do have my contact list saved with the Amazon Echo, but the people that I have on my contact list, a lot of them overlap among my family members. It will not be a big deal if somebody asks for a contact [from the smart speaker]."*

Nothing to hide from housemates: Several participants ($n = 6$), who shared their devices with their partners, believed that

it was vital for them to be able to check one another's information and that nothing should be hidden. P19 explicitly stated: *"... there should not be any secrets [between my wife and me]."* Those who shared the device with their friends emphasized the closeness they had with each other and explained that there was no need for them to withhold information from one another. P7, who shared his Amazon Echo Dot with his roommate, said: *"... there is just nothing to hide. We are pretty open with each other."*

Helplessness: Several participants ($n = 6$) expressed helpless acceptance of the privacy issues. They believed that there was no technology support that could help them. For example, P22 expressed concerns about his roommate gaining access to his calendar through Google Home, due to voice match false positives. He explained his lack of action as follows: *"... nothing ... there is nothing [that] can be done [to prevent my housemate from accessing my calendar]."*

Data control: Some participants believed that they had control over the data that they shared with other users. Specifically, two participants chose to link some information to the device and to intentionally hold back other information. To illustrate, P6 explained: *"... we are not sharing significant things, like some credit card information. I think it is up to you. You control what you share."* This statement indicates that the participant was concerned and unintentionally adopted a strategy of not sharing "significant things."

Strategies with Regard to External Entities

Participants reported adoption or awareness of diverse strategies for coping with their concerns about external entities.

Avoidance: Many participants ($n = 13$) intentionally avoided using certain features or storing sensitive information on the device. Similarly to many previous studies [1, 51], several of our participants ($n = 6$) chose not to use the purchasing feature of the device to prevent hackers from gaining access to their credit card information. Unique to our study is the finding that several participants ($n = 7$) avoided providing sensitive information to the speaker, including passwords, contact lists, calendars, and doctor's appointments. P16 illustrated her strategy as follows: *"From the very beginning, I did not [link] that much information [to the speaker]. So even if something goes wrong, a lot of my personal information is not as out there..."*

Sharing of responsibilities with the manufacturer: Few participants reported that they had taken every mitigation strategy that had been suggested to them or that they were aware of. Therefore, it was now the company's responsibility to ensure that the strategies users had taken effectively protected their data. P14 described the coping strategy that he adopted as follows: *"My Google account is pretty locked down. Like two-factor authentication ... and I have all of my security settings set up. ... So I think I am doing everything I can ... now it is up to [the speaker manufacturer]."*

Acceptance: Regarding external entities, many participants reported that they accepted the security and privacy risks identified, citing various reasons:

Trust in the company: Similarly to previous findings [56, 51], some participants believed that smart speaker manufacturers are trustworthy and that protecting users' privacy would be in the company's best interest [35]. Therefore, they trusted that the company would not misuse participants' data. In addition, we further explored how participants' **trust evolved based on news reports, both negative and positive**. Some participants recalled the news about Apple's refusal to mine data from an iPhone used by a terrorist [44] and expressed that they trusted smart speaker companies to also protect their data in a similar fashion. P12 stated his belief in Apple by saying, "[Apple] has always had a good business model and [business] ethics, [in] privacy as well. They have always been strong even with the government." At the same time, the attitudes of several participants were also influenced by negative news. Interestingly, participants intended to continue trusting the company as long they did not hear negative news about it, e.g.: "I have not heard anything bad happen [about the company]. If there was, then I would stop using [the speaker]" (P8).

Business model and technology improvement: Several (n = 7) participants believed that the company's collection of data was necessary to improve the technology. P7, who had a background in computer engineering, believed it was essential for the company to employ machine learning on a large amount of data to make its speaker "smarter." Others found it acceptable that the smart speaker companies used their data for generating revenue: "[The speaker companies] collect your data and use that to tailor ads to you. I think that is how they make money. So, I guess it makes sense for them to keep the data." (P8).

Helplessness: Several (n = 6) participants helplessly accepted the risks, expressing frustration with the current technology. For instance, when asked what he did to prevent the speaker company from collecting his data, P7 said: "Honestly, nothing. Because Google already has all the information it needs. It has everything that [is] digitally mine. There is nothing I can do about it." Other participants admitted to being concerned initially but said they eventually stopped being concerned when they realized that there was nothing they could do. As a result, the participants accepted the risk.

Data control: Participants also believed that they had control over the data they shared with the speaker manufacturers. On the one hand, participants believed they had the ability to choose the information that they linked to the device. P11 stated: "Anything that I do not need to or want to share, I just do not put it on [the speaker]. And the companies do not have access to that data." On the other hand, P21 believed that she could exercise control over her data at any time by unplugging the device: "I unplug the speaker so it cannot hear my conversation." These data control strategies suggest that participants were concerned in the first place and adopted corresponding strategies to address these concerns.

Comparative risk assessment: A couple of participants argued that their privacy could be violated through many other means, and there was no need to be specifically concerned about the smart speaker. For example, while explaining how credit card information can be stolen from smart speakers, P7 stated, "But [unauthorized access to financial information] can happen

with a credit card, either way. ... somebody can hack your bank account and steal money. So, the risks exist regardless."

Laziness: Although many participants identified security issues, they admitted not taking action to address their concerns. Some (n = 3) cited laziness as the reason for their behavior: "Human tendency is that you are always lazy. That is why [smart speakers] are popular." (P6).

DISCUSSION

Lessons Learned

Latent concerns. Based on their investigation of the link between privacy and security factors and the purchasing behaviors of IoT consumers, Emami-Naeini et al. [14] refer to the concerns that participants brought up after being prompted as "latent concerns." They suggest that latent concerns can surface readily for some consumers if privacy and security information is made salient.

Instead of prompting participants, we followed up on their responses of not being concerned and probed them for further explanations. We discovered that the adopted mitigation strategies and the concerns about using smart speakers were perceived by participants from two different perspectives, as Figure 3 illustrates. During the interviews, many participants first expressed their concerns about using a speaker. Then, they described their mitigation strategies (left rectangle in Figure 3). While other participants reported **not being concerned**, we found, however, that some of their justifications were actually unintentionally adopted strategies in the first place (right rectangle in Figure 3). For instance, several participants reported being OK with the speaker company collecting their data. They explained later that the reason for this acceptance was that they intentionally avoided "linking" sensitive information to the speaker. We consider such behavior (e.g., selecting which features to turn on) to be a coping strategy adopted in response to *latent concerns*. As employing avoidance strategy in response to latent concerns could deprive users from fully benefiting from their devices, corresponding solutions should be proposed to improve users' experiences.

Figure 3. Users' concerns, strategies, and justifications.

Users have concerns toward both external and internal entities. Unlike previous studies [54, 35], our findings suggest that users of shared speakers are concerned with risks emanating from both external (device manufacturers, governments, and hackers) and internal (housemates and visitors) entities.

Users adopted coping strategies to manage the risks without giving up using the device. For instance, by avoiding certain features (e.g., voice purchasing), some participants mitigated their concerns associated with these features, while enjoying others (e.g., playing music). These newly discovered concerns

about housemates and visitors suggest a need for more effective risk management to address users' concerns toward both types of entities.

Participants' concerns with respect to their housemates are often related to their mental models of speaker sharing. If participants with an advanced mental model have concerns about their housemates, they adopt corresponding strategies at the beginning, such as avoiding linking private information to the device. Since they had a good understanding of the underlying data sharing dynamics of the shared device, unlike participants with limited mental models, they were not worried about risks that could be easily prevented through device configuration. For example, impersonation as a concern identified by users with limited mental models, could be prevented by creating a personal voice profile. In contrast, participants with an advanced mental model have more concerns due to the current limitations of the technology, such as false positives in voice recognition. We did not find differences between primary and secondary users when it came to privacy concerns and risk management strategies.

Participants' perceptions of smart speaker risks are influenced by social media. As social media is becoming a major source of information for many people, most ($n = 23$) participants reported getting the news related to smart speakers through social networking (Facebook), micro-blogging (Twitter), photo-sharing apps (Instagram), and even subway TVs.

On the one hand, many participants tended to have greater trust that the company would protect their data if they had read positive news about the company's data sharing practices/policies. For instance, a few participants referred to the news that Apple refused to reveal customers' data to the FBI [15] to justify their trust in the company and its products. They were convinced by the news that Apple and other major companies would not reveal their data to the government and similar entities.

On the other hand, participants' concerns about the speaker tended to intensify if they saw negative news. Lau et al. [35] report that smart speaker *users* trust the device manufacturers (who also act as service providers for their devices) not to share user information with third parties. Some of our participants, however, suspected that these manufacturers could sell user data for profit, invading customers' privacy. Many of the suspecting participants referred to the Facebook data privacy scandal [7] to explain their doubts about Google and Amazon. The news led them to question whether the speaker manufacturer would indeed protect their data. Other negative news regarding device performance also affected participants' belief in the manufacturers' ability to secure their data.

The all-or-nothing (i.e., acceptance or avoidance) coping strategies suggest a lack of effective risk management. Although participants expressed different privacy concerns in regard to their housemates and third parties, they adopted similar coping strategies. On the one hand, participants commonly avoided using certain features to manage perceived risks. For example, many of them decided not to use the purchasing feature of the speaker to prevent either their housemates from unauthorized online shopping or hackers from gaining access

to their credit card information. However, such avoidance prevents users from making the full use of the speaker's capabilities, possibly reducing the perceived value of the technology. On the other hand, several participants expressed helpless acceptance of the perceived privacy and security risks. The employment of these avoidance and acceptance strategies suggests that technical support for effective risk management for smart speakers is yet to be developed by manufacturers and effectively employed by end users.

Multi-user features do not support users' needs well. Most speaker companies provide features, such as Amazon's Household profile and Google's voice match technology, that allow multiple users to share one device. However, these features were not widely adopted and appreciated by our participants. One possible reason is that many default settings accompany these features. For example, anyone in the household can access the contact list of the first user who links their Google account to the speaker if nobody's voice profile is set up [25], which caused privacy concerns among the participants. Additionally, when one user's voice profile is linked with the Google speaker, other users cannot use features such as the calendar and shopping lists if their voice profiles are not set up. It is not surprising that many participants voiced their dissatisfaction with this limitation.

A confounding factor was the lack of adequate mental models of how smart speakers were shared, no matter whether users' voice profiles were set up or not. Users' misunderstandings of how and what data are shared within the household could result in concerns with respect to their housemates.

Recommendations

Improve users' mental models. Improved user experience could help users with limited mental models to better understand their shared devices. For instance, when the household profile or voice match is first set up, the speaker companies can help users with limited mental models to learn (e.g., through tutorials) how the data will be shared or kept private. Therefore, users would not only understand the benefit that the features can provide, but also learn more about the privacy of the data.

Users should be made aware of how technology supports better risk management. As discussed above, many participants managed perceived risks by completely embracing or avoiding certain features. Some expressed their frustration about such an all-or-nothing trade-off between utility and privacy.

Furthermore, other participants admitted to being unaware of the full spectrum of mitigation strategies. They appeared to have developed a sense of helplessness with respect to the current technology. Thus, we believe that users of smart speakers should be helped to develop adequate mental models of the technology that will help them to make optimal risk management decisions. Our results show that when users have adequate mental models of their devices, they are able to make more effective risk management decisions. Users could also be given more options on how the data collected about them is managed, e.g., whether their voice data is collected, who the manufacture shares their data with, and how long the collected data is retained. Moreover, as suggested by [1],

this information should be delivered carefully to avoid being a burden on users. Both technology support for effective risk management and the development of adequate mental models through user experience are open research problems.

Trust in the device manufacturers needs to be developed and maintained. While data collection by such companies is a known concern of smart device users [41, 17, 56], we further discovered that participants raised concerns about the *purpose* and *scope* of the data collection. We suggest that companies present precise information to the public in a clear way. Instead of accusing the company's terms of service of being not clear enough to understand [1], a few of our participants said that reading the company's terms of service was tedious and onerous, which made them believe they were being tricked into accepting the terms. These findings indicate participants' lack of understanding of the data being collected by such companies, which led to many of their privacy concerns. Speaker manufacturers should build trust with their customers by, for example, helping the public to better understand their data collection, retention, and sharing practices (and possibly even the corresponding security controls, policies, and processes employed by the company).

Minimize trade-offs between convenience and privacy. Smart speaker companies should also explain and minimize the convenience-privacy trade-offs that they offer to the users. For instance, by turning on the "personal results," users give Google assistant access to 9 types of their data, including their calendar, contacts, and email [23]. Google should explain why a user who simply wants a personalized music playlist still needs to provide access to so much data. Since device brand is a factor that influences users' device adoption [14], we believe that maintaining such transparency will also benefit the speaker manufacturers in the long term.

Personalize information sharing for shared devices. Technology support for personalized sharing could facilitate more effective risk mitigation. To ease the concerns (including latent concerns) about other housemates accessing information through the shared smart speaker, users should be able to customize *what personal information* they are willing to share and the *people* they want to share with. For instance, users could selectively share their contact lists with their partners and intentionally withhold this information from their young children. After obtaining users' data sharing preferences, machine learning might be employed to predict such preferences. In the smart home domain, the preference results could be applied to other shared smart devices by users' choice, to improve their experiences. Thereby, users could choose to apply their data sharing decisions on all the shared devices.

Enhancing voice recognition technology. Improving voice recognition will alleviate risks emanating from housemates and visitors. Our results suggest that many participants were concerned that their smart speakers could not always distinguish users who had set up voice profiles. This was perceived as a risk because other users could misuse the device (e.g., for unauthorized purchases) if they could impersonate an authorized user. While an extra authentication mechanism can provide users with additional security and privacy, it could

also introduce more vulnerabilities. For example, Amazon enables users to add a 4-digit voice code to protect voice purchasing [5]. However, the user has to speak this code aloud for the speaker to process it, raising the possibility of the code being overheard by others. The same code is also stored in the Alexa App, which can be easily accessed if the phone is unlocked or through shoulder surfing. As voice recognition technology continues to mature [52], device manufacturers should improve the voice recognition on smart speakers to reduce false positive rates to acceptable levels.

LIMITATIONS

Since we aimed to study users' security and privacy concerns about their housemates, we recruited participants who shared their devices with their parents, siblings, children, partners, and roommates. However, there might be other types of relationships among users who share a household. Additionally, despite our efforts to study many types of smart speakers, we were unable to recruit any participants who own a smart speaker with a display. This was likely due to the relatively low market share of screen-equipped smart speakers at the time of our study [50]. The new capabilities of smart speakers with displays may raise unique concerns among users.

CONCLUSION

Smart speakers are rapidly gaining popularity. Despite the convenience offered by these devices, users have many security and privacy concerns. Our study investigated users' concerns with regard to their shared smart speakers and external entities. We explored the differences in users' corresponding coping strategies. Our findings reveal that participants were concerned primarily with their housemates' inappropriate access to personal information, and the misuse of the device by unintended users. We corroborate findings reported by others that, among the perceived risks regarding external entities, the use of the data being collected by the company is the major concern. Despite expressing different concerns about external entities versus housemates, our participants adopted similar risk management strategies for both types of adversaries. The adoption of all-or-nothing strategies suggests the lack of effective risk management by the users. We therefore offer several recommendations for future smart speaker design that might enable users to better cope with perceived risks. First, the design should allow users to personalize their sharing preferences. Second, the voice recognition technology should be improved to reduce false positives to acceptable levels. Finally, optimal risk management methods should be effectively communicated to end users.

ACKNOWLEDGEMENTS

This research has been supported in part by a gift from Scotiabank to UBC. We would like to thank members of the Laboratory for Education and Research in Secure Systems Engineering who provided their feedback on the reported research. Our anonymous reviewers provided important feedback and suggestions to improve the paper. Stylistic and copy editing by Eva van Emden helped to improve readability of this paper.

REFERENCES

- [1] Noura Abdi, Kopo M. Ramokapane, and Jose M. Such. 2019. More Than Smart Speakers: Security and Privacy Perceptions of Smart Home Personal Assistants. In *Proceedings of the Fifteenth USENIX Conference on Usable Privacy and Security (SOUPS'19)*. USENIX Association, Berkeley, CA, USA, 451–466. <http://dl.acm.org/citation.cfm?id=3361476.3361510>
- [2] Noura Aleisa and Karen Renaud. 2017. Privacy of the Internet of Things: a systematic literature review. In *Proceedings of the 50th Hawaii International Conference on System Sciences*.
- [3] Abdulaziz Alhadlaq, Jun Tang, Marwan Almaymoni, and Aleksandra Korolova. 2017. Privacy in the Amazon Alexa Skills Ecosystem. In *10th Workshop on Hot Topics in Privacy Enhancing Technologies (HotPETs)*.
- [4] Amazon. 2019a. Alexa Skills Kit. (2019). <https://developer.amazon.com/en-US/alexa/alexa-skills-kit>.
- [5] Amazon. 2019b. Require a Voice Code for Purchases with Alexa. (2019). <https://www.amazon.ca/gp/help/customer/display.html?nodeId=GAA2RYUEDNT5ZSNK>.
- [6] Amazon. 2019c. Set Up Contacts on Alexa App. (2019). <https://www.amazonforum.com/forums/devices/echo-alexa/502091-set-up-contacts-on-alexa-app>.
- [7] Mae Anderson. 2019. Facebook privacy scandal explained. (6 April 2019). <https://www.ctvnews.ca/sci-tech/facebook-privacy-scandal-explained-1.3874533>.
- [8] Joseph Bugeja, Andreas Jacobsson, and Paul Davidsson. 2016. On privacy and security challenges in smart connected homes. In *Intelligence and Security Informatics Conference (EISIC), 2016 European*. IEEE, Uppsala, Sweden, 172–175.
- [9] Nicholas Carlini, Pratyush Mishra, Tavish Vaidya, Yuankai Zhang, Micah Sherr, Clay Shields, David Wagner, and Wenchao Zhou. 2016. Hidden Voice Commands. In *USENIX Security Symposium*. USENIX Association, Berkeley, CA, USA, 513–530.
- [10] Hyunji Chung and Sangjin Lee. 2018. Intelligent Virtual Assistant knows Your Life. *CoRR* abs/1803.00466 (2018). <http://arxiv.org/abs/1803.00466>
- [11] Jamie Court. 2017. Google, Amazon Patent Filings Reveal Digital Home Assistant Privacy Problems. <https://www.consumerwatchdog.org/sites/default/files/2017-12/Digital%20Assistants%20and%20Privacy.pdf>. (2017). Last accessed: December 22, 2017.
- [12] Kiran K Edara. 2014. Key word determinations from voice data. (Aug. 5 2014). US Patent 8,798,995.
- [13] Pardis Emami-Naeini, Sruti Bhagavatula, Hana Habib, Martin Degeling, Lujo Bauer, Lorrie Faith Cranor, and Norman Sadeh. 2017. Privacy Expectations and Preferences in an IoT World. In *Proceedings of the Thirteenth USENIX Conference on Usable Privacy and Security (SOUPS '17)*. USENIX Association, USA, 399–412.
- [14] 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, New York, NY, USA, Article 534, 12 pages. DOI: <http://dx.doi.org/10.1145/3290605.3300764>
- [15] Chris Foxx and Dave Lee. 2016. Apple rejects order to unlock gunman's phone. (17 Feb. 2016). <https://www.bbc.com/news/technology-35594245>.
- [16] Batya Friedman, David Hurley, Daniel C. Howe, Helen Nissenbaum, and Edward Felten. 2002. Users' Conceptions of Risks and Harms on the Web: A Comparative Study. In *CHI '02 Extended Abstracts on Human Factors in Computing Systems (CHI EA '02)*. ACM, New York, NY, USA, 614–615. DOI: <http://dx.doi.org/10.1145/506443.506510>
- [17] Nathaniel Fruchter and Ilaria Liccardi. 2018. Consumer Attitudes Towards Privacy and Security in Home Assistants. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems (CHI EA '18)*. ACM, New York, NY, USA, Article LBW050, 6 pages. DOI: <http://dx.doi.org/10.1145/3170427.3188448>
- [18] Chuhan Gao, Varun Chandrasekaran, Kassem Fawaz, and Suman Banerjee. 2018. Traversing the Quagmire That is Privacy in Your Smart Home. In *Proceedings of the 2018 Workshop on IoT Security and Privacy (IoT S&P '18)*. ACM, New York, NY, USA, 22–28. DOI: <http://dx.doi.org/10.1145/3229565.3229573>
- [19] 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, New York, NY, USA, Article 268, 13 pages. DOI: <http://dx.doi.org/10.1145/3290605.3300498>
- [20] Dimitris Geneiatakis, Ioannis Kounelis, Ricardo Neisse, Igor Nai-Fovino, Gary Steri, and Gianmarco Baldini. 2017. Security and privacy issues for an IoT based smart home. In *Information and Communication Technology, Electronics and Microelectronics (MIPRO), 2017 40th International Convention on*. IEEE, Opatija, Croatia, 1292–1297. DOI: <http://dx.doi.org/10.23919/MIPRO.2017.7973622>
- [21] Barney G Glaser. 1978. Advances in the methodology of grounded theory: Theoretical sensitivity. (1978).
- [22] Stephanie Glen. 2016. Inter-rater Reliability IRR: Definition, Calculation. (2016). <https://www.statisticshowto.datasciencecentral.com/inter-rater-reliability/>.
- [23] Google. 2019a. Allow personal results on your shared Assistant devices. (2019). <https://support.google.com/assistant/answer/7684543>.

- [24] Google. 2019b. Link your voice to your Google Assistant device with Voice Match. (2019). <https://support.google.com/assistant/answer/9071681?hl.>
- [25] Google. 2019c. Make calls on Google Home. <https://goo.gl/uet3m2>. (2019).
- [26] Weijia He, Maximilian Golla, Roshni Padhi, Jordan Ofek, Markus Dürmuth, Earlene Fernandes, and Blase Ur. 2018. Rethinking Access Control and Authentication for the Home Internet of Things (IoT). In *Proceedings of the 27th USENIX Conference on Security Symposium (SEC'18)*. USENIX Association, Berkeley, CA, USA, 255–272.
- [27] Gary Horcher. 2019. Woman says her Amazon device recorded private conversation, sent it out to random contact. (2019). <https://www.kiro7.com/news/local/woman-says-her-amazon-device-recorded-private-conversation-sent-it-out-to-random-contact/755507974/>.
- [28] Catherine Jackson and Angela Orebaugh. 2018. A study of security and privacy issues associated with the Amazon Echo. *International Journal of Internet of Things and Cyber-Assurance* 1, 1 (2018), 91–100.
- [29] David Jonassen and Young Hoan Cho. 2008. *Externalizing Mental Models with Mindtools*. Springer, Boston, MA, 145–159. DOI: http://dx.doi.org/10.1007/978-0-387-76898-4_7
- [30] 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 USENIX Conference on Usable Privacy and Security (SOUPS '15)*. USENIX Association, Berkeley, CA, USA, 39–52. <http://dl.acm.org/citation.cfm?id=3235866.3235870>
- [31] Bret Kinsella. 2018. Amazon Alexa Now Has 50,000 Skills Worldwide, works with 20,000 Devices, Used by 3,500 Brands. <https://voicebot.ai/2018/09/02/amazon-alexa-now-has-50000-skills-worldwide-is-on-20000-devices-used-by-3500-brands/>. (2018). Last accessed: September 2, 2018.
- [32] Bret Kinsella. 2019. Loup Ventures Says 75% of U.S. Households Will Have Smart Speakers by 2025, Google to Surpass Amazon in Market Share. <https://voicebot.ai/2019/06/18/loup-ventures-says-75-of-u-s-households-will-have-smart-speakers-by-2025-google-to-surpass-amazon-in-market-share/>. (2019). Last accessed: June 18, 2019.
- [33] Bret Kinsella and Ava Mutchler. 2019. Smart Speaker Consumer Adoption Report 2019. https://voicebot.ai/wp-content/uploads/2019/03/smart_speaker_consumer_adoption_report_2019.pdf. (2019). Last accessed: March 20, 2019.
- [34] Todd Kulesza, Simone Stumpf, Margaret Burnett, and Irwin Kwan. 2012. Tell Me More?: The Effects of Mental Model Soundness on Personalizing an Intelligent Agent. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, New York, NY, USA, 1–10. DOI: <http://dx.doi.org/10.1145/2207676.2207678>
- [35] Josephine Lau, Benjamin Zimmerman, and Florian Schaub. 2018. Alexa, Are You Listening?: Privacy Perceptions, Concerns and Privacy-seeking Behaviors with Smart Speakers. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 102 (Nov. 2018), 31 pages. DOI: <http://dx.doi.org/10.1145/3274371>
- [36] Hosub Lee and Alfred Kobsa. 2016. Understanding user privacy in Internet of Things environments. In *Internet of Things (WF-IoT), 2016 IEEE 3rd World Forum on*. IEEE, Reston, VA, 407–412. DOI: <http://dx.doi.org/10.1109/WF-IoT.2016.7845392>
- [37] Hosub Lee and Alfred Kobsa. 2017. Privacy preference modeling and prediction in a simulated campuswide IoT environment. In *Pervasive Computing and Communications (PerCom), 2017 IEEE International Conference on*. IEEE, Kona, HI, 276–285.
- [38] X. Lei, G. Tu, A. X. Liu, C. Li, and T. Xie. 2018. The Insecurity of Home Digital Voice Assistants - Vulnerabilities, Attacks and Countermeasures. In *2018 IEEE Conference on Communications and Network Security (CNS)*. 1–9. DOI: <http://dx.doi.org/10.1109/CNS.2018.8433167>
- [39] Irene Lopatovska, Katrina Rink, Ian Knight, Kieran Raines, Kevin Cosenza, Harriet Williams, Perachya Sorsche, David Hirsch, Qi Li, and Adrianna Martinez. 2019. Talk to me: Exploring user interactions with the Amazon Alexa. *Journal of Librarianship and Information Science* 51, 4 (2019), 984–997.
- [40] Nathan Malkin, Joe Deatrick, Allen Tong, Primal Wijesekera, Serge Egelman, and David Wagner. 2019. Privacy Attitudes of Smart Speaker Users. *Proceedings on Privacy Enhancing Technologies* 2019, 4 (2019), 250–271.
- [41] Lydia Manikonda, Aditya Deotale, and Subbarao Kambhampati. 2018. What's Up with Privacy?: User Preferences and Privacy Concerns in Intelligent Personal Assistants. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society (AIES '18)*. ACM, New York, NY, USA, 229–235. DOI: <http://dx.doi.org/10.1145/3278721.3278773>
- [42] Benjamin Mayo. 2019. HomePod multi-user voice support and music handoff coming 'later this fall', new Ambient Sounds feature. <https://9to5mac.com/2019/09/11/homepod-features-later/>. (2019). Last accessed: December 25, 2019.
- [43] Terence V McCann and Eileen Clark. 2003. Grounded theory in nursing research: Part 1-Methodology. *Nurse Researcher (through 2013)* 11, 2 (2003), 7.

- [44] Ellen Nakashima. 2016. Apple vows to resist FBI demand to crack iPhone linked to San Bernardino attacks. (17 Feb. 2016). https://www.washingtonpost.com/world/national-security/us-wants-apple-to-help-unlock-iphone-used-by-san-bernardino-shooter/2016/02/16/69b903ee-d4d9-11e5-9823-02b905009f99_story.html.
- [45] Danny Palmer. 2018. Amazon's Alexa could be tricked into snooping on users, say security researchers. <https://www.zdnet.com/article/amazons-alexa-could-be-tricked-into-snooping-on-users-say-security-researchers/>. (2018). Last accessed: April 25, 2018.
- [46] Fahimeh Raja, Kirstie Hawkey, and Konstantin Beznosov. 2009. Revealing Hidden Context: Improving Mental Models of Personal Firewall Users. In *Proceedings of the 5th Symposium on Usable Privacy and Security (SOUPS '09)*. ACM, New York, NY, USA, 1:1–1:12. DOI: <http://dx.doi.org/10.1145/1572532.1572534>
- [47] Nirupam Roy, Haitham Hassanieh, and Romit Roy Choudhury. 2018a. BackDoor: Sounds that a microphone can record, but that humans can't hear. *GetMobile: Mobile Computing and Communications* 21, 4 (2018), 25–29.
- [48] Nirupam Roy, Sheng Shen, Haitham Hassanieh, and Romit Roy Choudhury. 2018b. Inaudible Voice Commands: The Long-Range Attack and Defense. In *15th USENIX Symposium on Networked Systems Design and Implementation (NSDI 18)*. USENIX Association, Renton, WA, 547–560. <https://www.usenix.org/conference/nsdi18/presentation/roy>
- [49] Paul Smith. 2018. Smart speakers and connected appliances the gateway drug as IoT goes mainstream. <https://www.afr.com/technology/gadgets/home-entertainment/smart-speakers-and-connected-appliances-the-gateway-drug-as-iot-goes-mainstream-20180513-h100i3>. (2018). Last accessed: May 14, 2018.
- [50] Greg Sterling. 2019. Alexa devices maintain 70% market share in U.S. according to survey. (9 Aug. 2019). <https://marketingland.com/alexa-devices-maintain-70-market-share-in-u-s-according-to-survey-265180>.
- [51] Madiha Tabassum, Tomasz Kosiński, and Heather Richter Lipford. 2019. "I Don't Own the Data": End User Perceptions of Smart Home Device Data Practices and Risks. In *Proceedings of the Fifteenth USENIX Conference on Usable Privacy and Security (SOUPS'19)*. USENIX Association, Berkeley, CA, USA, 435–450.
- [52] Naomi van der Velde. 2019. Innovative Uses of Speech Recognition Today. <https://www.globalme.net/blog/new-technology-in-speech-recognition>. (2019). Last accessed: July 08, 2019.
- [53] Yaxing Yao, Justin Reed Basdeo, Smirity Kaushik, and Yang Wang. 2019. Defending My Castle: A Co-Design Study of Privacy Mechanisms for Smart Homes. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM, New York, NY, USA, 198:1–198:12. DOI: <http://dx.doi.org/10.1145/3290605.3300428>
- [54] Eric Zeng, Shrirang Mare, and Franziska Roesner. 2017. End User Security & Privacy Concerns with Smart Homes. In *Proceedings of the Thirteenth USENIX Conference on Usable Privacy and Security (SOUPS '17)*. USENIX Association, Berkeley, CA, USA, 65–80. <http://dl.acm.org/citation.cfm?id=3235924.3235931>
- [55] Guoming Zhang, Chen Yan, Xiaoyu Ji, Tianchen Zhang, Taimin Zhang, and Wenyan Xu. DolphinAttack: Inaudible Voice Commands. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (CCS '17)*. ACM, New York, NY, USA, 103–117. DOI: <http://dx.doi.org/10.1145/3133956.3134052>
- [56] Serena Zheng, Noah Apthorpe, Marshini Chetty, and Nick Feamster. 2018. User Perceptions of Smart Home IoT Privacy. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 200 (Nov. 2018), 20 pages. DOI: <http://dx.doi.org/10.1145/3274469>