

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/335091754>

Alexa's Voice Recording Behavior: A Survey of User Understanding and Awareness

Conference Paper · August 2019

DOI: 10.1145/3339252.3340330

CITATIONS

23

READS

2,165

3 authors, including:



[Yousra Javed](#)

Illinois State University

31 PUBLICATIONS 248 CITATIONS

SEE PROFILE

Alexa's Voice Recording Behavior: A Survey of User Understanding and Awareness

Yousra Javed
yjaved@ilstu.edu
Illinois State University
IL, USA

Shashank Sethi
ssethi5@ilstu.edu
Illinois State University
IL, USA

Akshay Jadoun
asjado1@ilstu.edu
Illinois State University
IL, USA

ABSTRACT

The use of Amazon's virtual assistant *Alexa* in controlling smart home devices is on the rise. The convenience provided by an Alexa-enabled device comes at the cost of Alexa service's voice recording and storage behavior, raising privacy concerns. Amazon claims to record and store voice data in the cloud only when the wake word is spoken. However, Alexa records user's voice even at times when the wake word is not used. Though short, these recordings can potentially contain a sensitive conversation between individuals. Anyone in the possession of the Amazon Alexa mobile application with which the smart device is registered, can access these recordings.

This paper aims to investigate Alexa users' understanding and awareness of its voice recording and storage behavior. We conducted an Amazon Turk study on 113 Amazon Echo owners. Our results show that 91% of the participants had at-least one instance of unintended voice recording. 53.09% of our participants reported that a subset of these unlabeled recordings did not contain any voice commands. Although a majority of the participants (64.6%) were aware that Alexa stores data in Amazon cloud, the participants lacked awareness of 1) intended and unintended voice data being recorded and stored, 2) entities who can potentially access these voice recordings, and 3) voice recording deletion. Having technical/CS background did not influence the participants responses. However, the participants who were strongly concerned about privacy (fundamentalists) seemed to have a better understanding of Alexa's voice recording storage behavior.

KEYWORDS

Amazon Alexa, IoT security and privacy, voice recording, Internet-connected devices

ACM Reference Format:

Yousra Javed, Shashank Sethi, and Akshay Jadoun. 2019. Alexa's Voice Recording Behavior: A Survey of User Understanding and Awareness. In *Proceedings of the 14th International Conference on Availability, Reliability and Security (ARES '19)* (ARES '19), August 26–29, 2019, Canterbury, United Kingdom. ACM, Canterbury, UK, 10 pages. <https://doi.org/10.1145/3339252.3340330>

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

ARES '19, August 26–29, 2019, Canterbury, United Kingdom

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-7164-3/19/08...\$15.00

<https://doi.org/10.1145/3339252.3340330>

1 INTRODUCTION

Virtual assistants have gained popularity over the recent years. These artificial intelligence companions are playing an important role in our daily life by performing tasks such as checking weather, playing music, or setting reminders etc in response to simple voice commands. Some of the popular virtual assistants include Amazon's Alexa, Apple's Siri, Google's Assistant, Microsoft's Cortana, and Samsung's Bixby. Alexa has integration with third-party applications like uber [16] and unlike other virtual assistants contains around 15000+ skills [11].

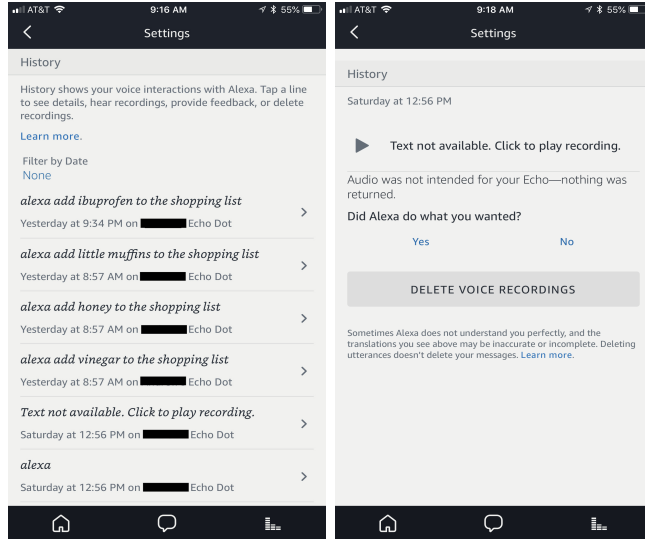
With the growing adoption of internet of things (IoT), virtual assistants are now being embedded in home devices. These assistants are being used to control items around the house—from light bulbs to coffee makers—all with a simple voice command. One such virtual assistant-enabled device is a smart speaker. Popular smart speakers include Amazon's Alexa powered Echo [5, 20].

Smart speakers and other speech-activated devices provide several convenience and benefits. However, they also raise numerous privacy concerns. In 2015, Samsung's microphone-enabled SmartTV was reported to be "always on" in violation of federal wiretapping laws. Samsung's Privacy Policy warned that sensitive conversations might be swept up and transmitted to third-parties as part of the TV's voice controlled search function. Amazon Alexa's design emphasizes user awareness, consent-based features, user control over the device, and incorporates the following emerging privacy-conscious practices stated in [6]:

1. Amazon's terms of use clearly state that Alexa stores voice commands as text and voice recording on Amazon cloud. These recordings are used to improve the quality of Alexa. Amazon claims that it does not collect or disclose recorded voice data (or text translations) beyond what the user reasonably expects by stating that the voice is only recorded when the wake word is used to activate Alexa [3].
2. Alexa-enabled device provides visual cues to indicate voice recording/transmission.
3. The hard on/off switch on the Alexa-enabled device can be used to disable the microphone to give user control over the device.
4. The device owner can view these recordings from their Alexa mobile application and has the option to delete them [4].

We observed that there are instances when Alexa records voice even when the wake word is not spoken. These instances can occur due to virtual assistant's misinterpretation of a spoken word with its wake word and could potentially contain sensitive information. The unintended voice recordings can be found in the history section inside Alexa application's settings tab and are titled "*Text not available. Click to play recording*" (see Figure 1(a)). Thus, any person who has access to the smart device owner's Alexa mobile application can

view/listen to the previous voice commands given to Alexa and additionally, the unintended voice recordings. Alexa's terms of use do not mention unintended voice recording. Moreover, no visual cues are provided to indicate unintended voice recording/transmission.



(a) History containing an instance of un-labeled and potentially unintended voice recordings (2nd row from bottom)

Figure 1: Voice command history inside Alexa application

The occurrence of these unintended voice recordings motivated us to investigate Alexa users' understanding and awareness of the virtual assistant's voice recording and storage behavior. More specifically, we are interested in answering the following research questions:

- RQ1 Are Alexa users aware that their voice commands are stored on Amazon cloud?
- RQ2 Do Alexa users understand who can access their voice recordings and how to delete them?
- RQ3 Do Alexa users have unintended voice recording instances where Alexa records voice in the background and stores them as commands?
- RQ4 Do Alexa users consider Alexa's unintended voice recording behavior a privacy concern?

Existing security and privacy research on virtual assistants and other microphone-enabled devices have proposed potential unintended voice recording as a privacy and security concern [7]. McRenyolds et al. have interviewed parents and children to understand user mental models and privacy concerns for internet-connected toys [15]. However, to the best of our knowledge no work investigates user awareness and concerns of Alexa's unintended voice recording.

We conduct a user study on the crowdsourcing website, Amazon Mechanical Turk ¹ to investigate our research questions. Our results show that although a majority of the participants (64.6%)

¹<https://www.mturk.com>

were aware that Alexa stores data in Amazon cloud, the participants lacked awareness of 1) intended and unintended voice data being recorded and stored, 2) entities who can potentially access these voice recordings, and 3) voice recording deletion. Having technical/CS background did not influence the participant responses. However, the participants who were strongly concerned about privacy (fundamentalists) seemed to have a better understanding of Alexa's voice recording storage and access behavior. 91% of the participants had at-least one instance of unintended voice recording. 52% of the participants considered Alexa's unintended voice recordings a privacy concern.

The rest of this paper is organized as follows. Section 2 describes Amazon Alexa's architecture and its terms of use on voice recording. Section 3 discusses the related work. Section 4 describes our methodology including the user study design and results. Section 5 discusses the results and presents the limitations of our work. Finally, Section 6 presents our conclusions.

2 AMAZON ALEXA

2.1 Architecture

The main components of Amazon Alexa's architecture can be grouped into two categories, i.e., the client side and the cloud side.

1. Client Side: The client side comprises of two endpoints:

- An Alexa-enabled device such as Amazon Echo.
- A companion application that needs to be installed on the user's mobile device. This application can be used to register smart home devices, add skills, and to view, listen, and delete voice command history.

2. Cloud Side: The cloud side consists of the virtual assistant Alexa that operates in Amazon's cloud environment. When the user says the wake word e.g., "Alexa", their voice command is recorded and sent for processing and storage to the Amazon cloud. Figure 2 shows the Amazon Alexa's architecture.

2.2 Voice Recordings

2.2.1 Storage. Amazon's *Terms, Warranties, and Notices* document [3] states that when a user speaks to Alexa, a recording of what they asked Alexa is sent to Amazon's cloud where the request is processed along with other information to respond to the user. For example, if a user asks "Alexa, play top hits on Amazon Music", the recording of the request is used along with the information from Amazon Music to play top hits. Amazon's terms also state that by default, Echo devices are designed to only detect the user's chosen wake word (Alexa, Amazon, Computer or Echo). The device detects the wake word by identifying acoustic patterns that match the wake word. No audio is stored or sent to the cloud unless the device detects the wake word (or Alexa is activated by pressing a button).

While the Alexa application might store certain data (e.g., map data to improve the performance) locally, it does not store any voice recordings.

2.2.2 Review and Deletion. Users can review the voice recordings associated with their account and delete these voice recordings one by one (See Figure 1(b)) or all at once by either visiting **Settings > Alexa Account > Alexa Privacy** in the Alexa application or by

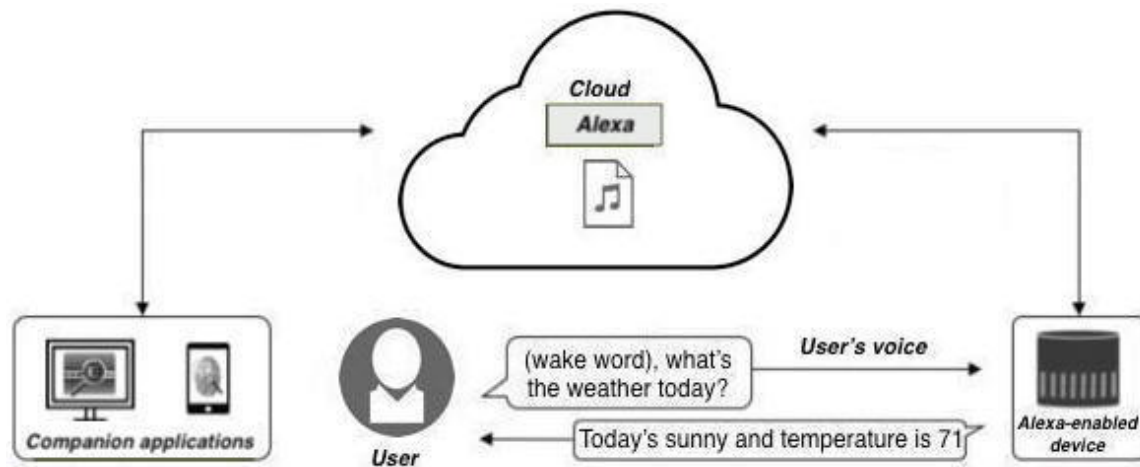


Figure 2: Amazon Alexa's Architecture

visiting Alexa's privacy website [4]. Users can also review these voice recordings by visiting **Settings > Alexa Account > History** in the Alexa application and delete them all at once for each of the Alexa-enabled products by visiting the website for managing Alexa content and devices [1].

2.3 Alexa Skills

Users can view and enable skills (voice-driven Alexa capabilities) in their Alexa application or by voice. When the user uses a skill, related information may be exchanged with the developer of that skill, such as user's answers when playing a trivia skill, user's ZIP code when asking for the weather, or the content of user's requests.

3 RELATED WORK

The literature most related to our work falls into four categories 1) Privacy concerns related to virtual assistants 2) User mental models and opinions of virtual assistants 3) MTurk surveys of user attitudes towards privacy and 4) Uses of Westin Privacy Index.

3.1 Privacy concerns related to virtual assistants

User interaction with virtual assistants constantly generates voice data that gets stored in the cloud. Gray [6] provides a comprehensive overview of the privacy implications of having always-on microphone-enabled devices. These include biometric identification and one and two party consent. The author also discusses privacy-conscious practices to alleviate these privacy concerns. These practices include 1) providing visual cues that clearly indicate when the device is recording and/or transmitting information, 2) use limitation, and 3) the ability to access and delete stored voice data. Pfeifle [17] states an additional threat to user privacy—that the law enforcement can use such devices in new ways that users are not prepared for during investigations.

Recently, Amazon introduced Alexa skills kit platform to enable developers to build voice-driven capabilities, called skills, for Alexa. AlHadlaq et al. [2] analyzed the privacy policies set up by the

developers for all Alexa skills available on Amazon. They found out that 75% of the skills lack a privacy policy that describes the use and distribution of customer's personal information.

3.2 User mental models and opinions of virtual assistants

McRenyolds et al. [15] have investigated user mental models and privacy expectations of Internet-connected toys such as "Hello Barbie". These toys are microphone-enabled and use speech recognition to respond to questions and answers. The authors' interviews with parents show that many parents voiced privacy concerns and were sensitive to the issues surrounding monitoring what their child does with the toy as well as what data would be recorded and retained by the company. The authors also interviewed the children and found out that the children were not aware that the toys were recording or that the recordings were accessible to their parents.

3.3 MTurk surveys of user attitudes towards privacy

Amazon Mechanical Turk (MTurk) is a crowdsourcing platform widely used to conduct behavioral research, including studies of online privacy and security. Lee et al. [13] surveyed 200 participants over MTurk on hypothetical IoT scenarios to identify contextual parameters that are associated with higher or lower acceptance of sensor tracking in IoT environments. Kang et al. [9] conducted an MTurk-based survey of how users manage their personal information online and what are their policy preferences about anonymity. Their results show that Indian MTurk workers are much less concerned than American workers about their privacy and more tolerant of government monitoring. We draw our sample from the United States. Kelley [10] conducted MTurk-based studies on user attitudes and perceptions related to phishing, targeted advertising, and privacy labels.

MTurk has also been used to investigate privacy issues related to social media. Wang et al. [19] investigated American, Chinese, and Indian social networking site (SNS) users' privacy attitudes and

practices using MTurk-based survey. Their analysis showed that American respondents were the most privacy concerned, followed by the Chinese and Indians. Liu et al. [14] conducted an MTurk-based study to investigate the disparity between Facebook users' desired and actual privacy settings and found that the privacy settings match users' expectations only 37% of the time.

3.4 Uses of Westin Privacy Index

Several studies have investigated the correlation between the Westin categories (fundamentalist, pragmatist, and unconcerned) and actual or intended behaviors. Woodruff et al. [21] conducted a study to understand the behavioral intentions and consequences of users with respect to their Westin privacy index. They found no correlation between the Westin categories and behavioral intent, as well as between the Westin categories and consequences. Similarly, Jensen et al. [8] investigated privacy practices of Internet users. They observed that Westin "privacy fundamentalists", do not appear to form a cohesive group for privacy-related decision making. On the other hand, a study on Snapchat showed that privacy unconcerned respondents are slightly more likely to send "photos/videos of themselves" (62.5%) than Pragmatists (31%) or Fundamentalists (28%) [18]. We are interested in investigating if there is any correlation between Westin categories and participant understanding and awareness of Alexa's voice recording behavior.

4 RESEARCH METHODOLOGY

To address our research questions regarding user understanding and awareness of Alexa's voice recording behavior, we formulate the following hypotheses:

- H1 Users lack comprehension of Alexa's voice recording storage, access, and deletion practices
- H2 Alexa records voice in the background even when no wake word is spoken
- H3 Users consider Alexa's unintended voice recording a privacy concern

4.1 Study Design

To evaluate our hypotheses, we designed a survey-based user study focused on the virtual assistant Alexa. Our survey (see Appendix) is comprised of the following sections:

1. Screening question: Since we required participants who owned an Alexa-enabled smart device, more specifically a smart speaker (Echo and its newer versions), we filtered participants who did not own such a device. To ensure that the participants owned an Alexa-enabled device, they were required to go to their Alexa application's settings and report the software version used.

2. Demographics and virtual assistant usage: The eligible participants first answered questions about their demographic background. We also asked the participants which other virtual assistants they used and how long they have used it.

3. Knowledge about virtual assistant's data recording storage, access, and deletion behavior: The next section of the survey asked the participants about their perceptions related to what

data (if any) is collected by the virtual assistant and where is this data stored. The participants also specified who has access to this data. Lastly, participants stated if they can delete the stored data and if so how.

4. Awareness and perceptions about virtual assistant's unintended voice recording: In this section, we asked the participants to go through the history tab from the settings menu of their Alexa application and locate recordings (within the last three months) that were titled "*Text not available. Click to play recording*". From here onwards, we refer to these recordings as unlabeled recordings. The three month time period was chosen to make it easier for participants to calculate the approximate number of unlabeled recordings. The participants were asked to report the approximate number of such recordings. To ensure that such instances existed in a participant's Alexa application, they were asked to provide a screenshot of one such instance (similar to Figure 1(a)) by clicking on it. These screenshots did not contain any visible private information. If the participants had instances of unlabeled recordings, they were instructed to listen to a few of them. The participants then reported on the percentage of such recordings that were not voice commands and the kind of information present in those recordings. Each participant expressed their level of surprise after they learnt about these recordings and whether they considered this a privacy issue.

5. Westin's privacy classification: Finally, the participants answered three questions (Q.11 in Appendix A.4) [12] to be classified into one of the three categories:

- (1) Fundamentalists - Strongly concerned about privacy
- (2) Pragmatists - Moderately concerned about privacy
- (3) Unconcerned - Unconcerned about privacy

The study was set up as a Human Intelligence Task (HIT) on Amazon Mechanical Turk and was approved by our institution's human subjects review board (IRB). We did not seek to elicit private information from the participants, and only asked questions about their perceptions and facts. To ensure that we receive meaningful data, we added attention-check questions in the survey, in addition to the Alexa-enabled device ownership verification questions. Moreover, in order to qualify for the HIT, the workers were required to have a minimum HIT acceptance rate of 96%. The average completion time for the task was 15-20 minutes. Each participant was paid \$0.50 for their time.

4.2 Results

4.2.1 Participant Demographics: A total of 150 participants from the United States completed the HIT between 1st September 2018 and 1st October 2018. After filtering the responses based on attention-check and device verification questions, the remaining number of responses were 113. A majority of the participants were male, aged between 20 and 30 years. More than 50% of our participants were white and had at least 4 years of college education. 52% of the participants did not have a Computer Science or technical background, and had used an Alexa-enabled device for less than a year. Table 1 shows our participant demographics.

		% of participants
Age	18-20	4.28
	20-30	45.71
	30-40	37.85
	40-50	7.14
	50-60	3.5
	60 and above	1.42
Gender	Male	60.71
	Female	39.28
Ethnicity	White	67.14
	Hispanic	7.85
	Asian	13.57
	Black	7.85
	Hawaiian	0.71
	Native	1.42
	Decline to answer	1.42
Education	High School	8.57
	2 years of college	18.57
	4 years of college	50.71
	> 4 years of college	20.71
	Decline to answer	1.42
CS Background	Yes	47.14
	No	52.14
Westin's Privacy Index	Fundamentalist	19.28
	Pragmatist	74.28
	Unconcerned	6.42
Duration of Alexa use	0 - 6 months	12.85
	6 months - 1 year	44.28
	1 year - 2 years	30.71
	2 year - 3 years	9.28
	More than 3 years	2.85

Table 1: Participant Demographics

We classified our participants based on Westin's privacy index using their responses to the Westin questions. 19.28% of our participants were fundamentalists, 74.28% were pragmatists, and 6.42% were unconcerned. Therefore, a majority of our participants (pragmatists) were moderately concerned about privacy.

Figure 3(a) shows the tasks for which our participants used Alexa. The top three tasks included setting reminders and alarms, getting information such as current weather, and listening to music and audiobooks.

4.2.2 Hypothesis 1: Our first hypothesis tests whether users have incorrect perceptions about Alexa's voice recording storage, access, and deletion practices. Table 2 summarizes the percentage of participants who correctly answered questions regarding Alexa's voice recording storage, access, and deletion.

Voice Recording Storage. We were interested in understanding what percentage of the participants are aware that their voice commands are being recorded and stored on Amazon cloud.

For this purpose, we first looked at the participant responses to the following question: *In your opinion, does Amazon Alexa store*

any data? If yes, what data? We coded the participants' responses manually. All responses containing words/sentences with meaning similar to *voice commands* were marked as correct responses. These responses included words such as *voice*, *recording(s)*, *command(s)*, and *conversation(s)*. Overall only 36.28% of the participants reported that Alexa stores their voice commands. 1-sample proportions test (binomial test) with continuity correction ($\chi^2 = 7.9646$, $df = 1$, $p\text{-value} = 0.0023$) showed that less than 50% of the participants had knowledge about Alexa's voice recording storage. Next, we analyzed if there is a difference in the perceptions of participants who have CS/technical background and those who do not. Only 34.14% of participants who had CS/technical background reported that Alexa stores voice commands, whereas 36.61% of participants who had no CS/technical background reported that Alexa stores voice commands. Pearson's Chi-squared test ($\chi^2 = 0.62237$, $df = 2$, $p\text{-value} = 0.7326$) showed no significant differences between the responses of the two groups. We also analyzed whether there is a difference between the perceptions of participants who are concerned about privacy and those who are not. For this purpose, we used participants' Westin's privacy index. 55.5% of fundamentalists, 44.4% of unconcerned, and 21.1% of pragmatists reported that Alexa stores voice commands. Pearson's Chi-squared test ($\chi^2 = 7.1906$, $df = 2$, $p\text{-value} = 0.02745$) showed a significant difference between the responses of the three groups. Pairwise comparisons showed the largest difference occurred between fundamentalists and pragmatists.

To understand if participants are aware of the location where the data recorded by Alexa is being stored, we analyzed participants' response for the location of the type of data (stored by Alexa) they reported earlier. Once again, we coded the participants responses manually. All responses containing words/sentences with meaning similar to *Amazon Cloud* were marked as correct responses. These responses included words such as *Amazon cloud*, *cloud*, *Amazon Server(s)*, *Amazon*. Overall, 64.6% of the participants reported that Alexa stores their data on Amazon cloud. Binomial test ($\chi^2 = 9.0619$, $df = 1$, $p\text{-value} = 0.00261$) showed that the more than 50% of the participants had knowledge about Alexa's voice recording storage location. 60.9% of participants who had CS/technical background reported that Alexa stores data in Amazon cloud, whereas 66.1% of participants with no CS/technical background reported that Alexa stores data in Amazon cloud. Pearson's Chi-squared test ($\chi^2 = 2.0783$, $df = 2$, $p\text{-value} = 0.3538$) showed no significant differences between the responses of the two groups. Similarly, 85.18% of the fundamentalists, 55.5% of unconcerned, and 43.26% of the pragmatists reported that Alexa stores data in Amazon cloud. Pearson's Chi-squared test ($\chi^2 = 6.6462$, $df = 2$, $p\text{-value} = 0.03604$) showed a significant difference between the responses of the three groups. Once again, the pairwise comparisons showed the largest difference occurred between fundamentalists and pragmatists.

Voice Recording Access. We investigated participants' understanding of who can access the data stored by Amazon Alexa. For this purpose, we looked at the participant responses to the question: *In your opinion, which persons/entities can potentially access this data?* Once again we coded the participant responses manually. Responses that included wording similar to *Amazon* or *anyone who has access to Alexa-enabled device owner's mobile application/Amazon*

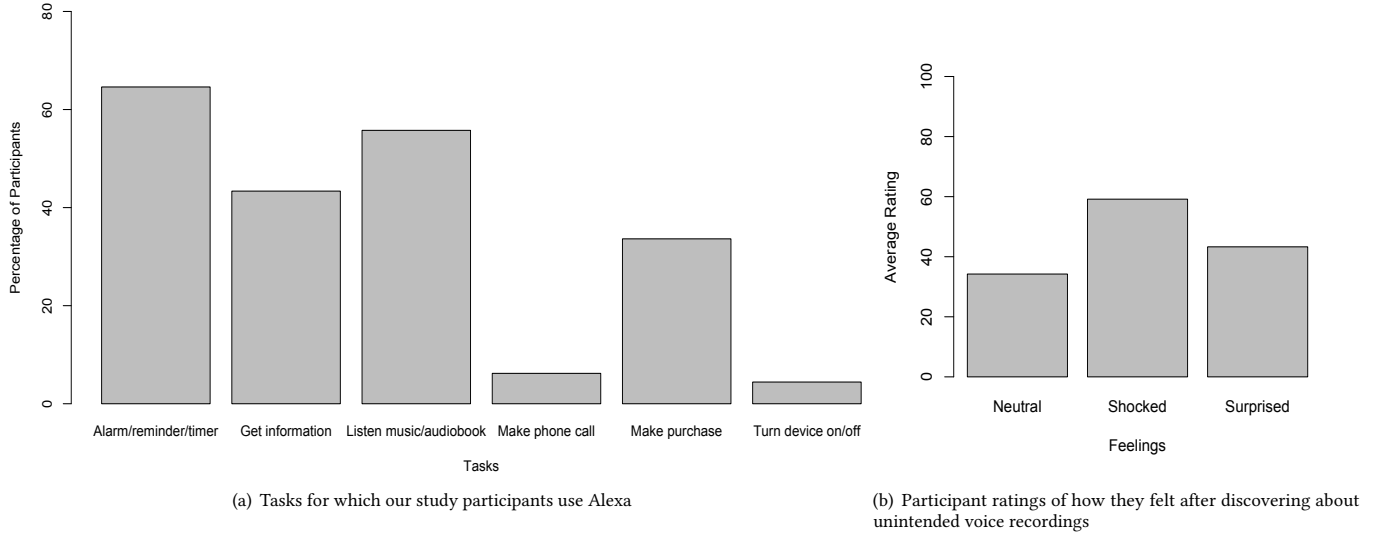


Figure 3: Participant responses for Alexa use and feelings towards unintended voice recordings

account were labeled as correct responses. Overall, 28.31% of the participants correctly answered this question. Binomial test ($\chi^2 = 20.389$, $df = 1$, $p\text{-value} = 3.159e^{-06}$) showed that the less than 50% of the participants had knowledge about entities who can access data stored by Amazon Alexa.

Analysis with respect to CS/technical background showed that 31.7% of the participants with CS/technical background and 26.7% of participants with no CS/technical background answered the data access question correctly. However, Pearson's Chi-squared test ($\chi^2 = 0.71191$, $df = 2$, $p\text{-value} = 0.7005$) showed no significant differences between the responses of the two groups. Analysis with respect to Westin's privacy index showed that 33.3% of the fundamentalists, 11.1% of unconcerned, and 25% of the pragmatists answered the question correctly. However, Pearson's Chi-squared test ($\chi^2 = 1.3211$, $df = 2$, $p\text{-value} = 0.5166$) showed no significant difference between the responses of the three groups.

Voice Recording Deletion. To assess participants' understanding regarding deletion of the data stored by Amazon Alexa, we looked at participant responses to the question: *In your opinion, can you delete the data stored by Amazon Alexa?* The possible responses were Yes, No, and I Don't Know, with the correct answer being "Yes". Overall, only 16.8% of the participants correctly answered this question. Binomial test ($\chi^2 = 48.46$, $df = 1$, $p\text{-value} = 1.685e^{-12}$) showed that less than 50% of the participants had knowledge that the data stored by Alexa can be deleted.

Analysis with respect to CS/technical background showed that 12.1% of participants who had CS/technical background correctly answered that they can delete the data. Whereas, 19.7% of participants with no CS/technical background answered the data deletion question correctly. However, Pearson's Chi-squared test ($\chi^2 = 1.2557$, $df = 2$, $p\text{-value} = 0.5337$) showed no significant differences between the responses of the two groups. Analysis with respect

to Westin privacy index showed that 18.5% of the fundamentalists, 37.5% of unconcerned, and 14.1% of the pragmatists answered the question correctly. However, Pearson's Chi-squared test ($\chi^2 = 2.9135$, $df = 2$, $p\text{-value} = 0.233$) showed no significant difference between the responses of the three groups.

We also asked the participants how they will delete the data retained by Alexa. The participant responses were coded manually. Responses with wording similar to *through the Alexa mobile application, or through the Amazon website* were marked as correct. Binomial test ($\chi^2 = 6.4796$, $df = 1$, $p\text{-value} = 0.005456$) showed that only 18.5% of the participants correctly answered the question. Similarly, 12.2% of participants with CS/technical background and 22.5% of the participants with no CS/technical background answered the question correctly with Pearson's Chi-squared test ($\chi^2 = 2.0669$, $df = 2$, $p\text{-value} = 0.3558$) showing no difference between the responses of the two groups. With regards to the privacy classification, 22.2% of fundamentalist, 33.3% of unconcerned, and 15.3% of pragmatists answered the question correctly. However, Pearson's Chi-squared test ($\chi^2 = 2.6558$, $df = 2$, $p\text{-value} = 0.265$) showed no significant difference between the responses of the three groups.

4.2.3 Hypothesis 2: Our second hypothesis investigates whether Alexa records voice in the background even when no command is given. These voice recordings are unlabeled and do not display the command given to Alexa in a text format. We were interested in investigating whether our participants had at least one such unlabeled voice recording in the three months (June 2018 - September 2018) prior to the time of participation. We requested the participants to report the approximate number of such recordings. 91% of the participants had at least one such instance. Binomial test ($\chi^2 = 74.903$, $df = 1$, $p\text{-value} < 2.2e^{-16}$) showed that more than 80% of the participants had at least 1 instance of unintended recording, with the average number of such instances per participant being 27.

		Storage		Access	Deletion	
		What data is stored?	Where is this data stored?	Who can access this data?	Can you delete this data?	How to delete this data?
Overall		36.28%	64.6%	28.31%	16.8%	18.5%
Educational background	CS	34.14%	60.9%	31.7%	12.1%	12.2%
	Non-CS	36.61%	66.1%	26.7%	19.7%	22.5%
Privacy Index	Fundamentalists	55.5%	85.18%	33.33%	18.5%	22.2%
	Pragmatists	21.1%	43.26%	25%	14.1%	15.3%
	Unconcerned	44.4%	55.5%	11.1%	37.5%	33.3%

Table 2: Percentage of participants who correctly answered questions regarding Alexa's voice recording storage, access, and deletion

Further, we investigated if these unlabeled recordings contained participants' misinterpreted voice commands/wake word or some other information. 53.09% of our participants reported that a subset of these unlabeled recordings did not contain any voice commands whereas 46.9% participants reported that the unlabeled recordings only contained misinterpreted voice commands/wake words. However, Binomial test ($\chi^2 = 0.31858$, $df = 1$, $p\text{-value} = 0.2862$) did not show any significant difference between the participants having unintended voice recording and those having legitimate but misinterpreted voice commands.

We asked the participants to report the type of information that was recorded by Alexa. Below are the three categories of information present in participant responses:

- Conversations over the phone or in-person
- Radio, TV sounds, or movie playing in the background
- Background noise

5.3% of our participants reported that all of their unlabeled recordings contained sensitive audio data, whereas 29.2% reported that some of their unlabeled recordings contained sensitive audio data.

4.2.4 Hypothesis 3: Our third hypothesis investigates whether participants considered Alexa's voice recording behavior a privacy concern. Participants were asked to report on a Likert scale whether they considered Alexa's voice recording behavior a privacy concern. 52% participants somewhat or strongly agreed to this Alexa behavior as a privacy concern. However, Binomial test ($\chi^2 = 0.14159$, $df = 1$, $p\text{-value} = 0.6466$) showed no significant difference between the proportion of participants who considered this behavior a privacy concern and those who did not.

Participants were also asked to report (on a scale of 0 to 100) on how they felt after discovering the unlabeled recordings containing information other than voice commands. Figure 3(b) shows the average ratings of participant feelings. The average participant rating for being neutral was 34.23 whereas, the average ratings for being surprised and shocked were 43.29 and 59.15 respectively. Therefore, participants did seem to be concerned after discovering unintended voice recordings. For example, one participant commented: *"There are 5 or 6 of my phone calls recorded, nothing really important but it surprised me"*. Other participants expressed their concern upon finding out that their private conversations were recorded: *"This is a little freaky. There were conversations with my wife, there was us talking about the neighbors dog and shooting it because it was*

barking, there was a recording of me snoring in the living room. a bunch of other kinda private stuff".

5 DISCUSSION

Voice Recording Storage, Access, and Deletion : A majority of our study participants (64.6%) were aware that Alexa stores data in Amazon cloud. Since cloud is a buzzword, many participants probably have heard of this term from others. However, the participants were not sure of the kinds of data (especially voice recordings) are being stored in the cloud. Only 36.28% participants reported that Alexa stores their *voice commands* in the cloud. Similarly, only 28.31% participants reported that Amazon and anyone else who has access to participants' Alexa mobile application can access their voice recordings. Lastly, only 16.8% participants reported that they can delete the voice recordings. Out of these participants, only 18.5% correctly answered how they can delete the data through the Alexa mobile application or through Amazon's website. Overall, participants had the least knowledge about data deletion. Therefore, they were unaware that they could check the history of their voice commands, and that there was an option to delete voice recordings from their Alexa application.

We observed that there was no correlation between technical/CS background and comprehension. However, the participants who were strongly concerned about privacy (fundamentalists) did seem to have a better understanding that voice recordings are being stored.

Unlabeled Voice Recordings : Our results showed that 91% of the participants had instances of unlabeled voice recordings. 53% of the participants reported that a subset of their unlabeled voice recordings did not contain voice commands and included a sensitive conversation. Therefore, contrary to its terms, Alexa does record voice when the wake word is not provided. However, we did not investigate the circumstances/scenarios when Alexa records voice in the background unbeknownst to the user. Participants were surprised and shocked to find out such instances from inside their Alexa application and 52% of the participants considered this behavior a privacy concern.

Limitations: Our study is not without limitations. We recruited participants from the United States. Therefore, our results are not a representative of Alexa users from all across the world. Moreover, we did not have an equal number of participants from all age groups,

genders, ethnic groups, educational backgrounds, and duration of Alexa use. These demographic factors could have impacted the results. For example, a majority of the participants were relatively new users of Alexa, which could have impacted their understanding and awareness.

6 CONCLUSION

We investigated user understanding and awareness of Alexa's voice recording behavior. Our results show that Alexa users have instances where Alexa recorded their voice without their intention. Users are also concerned about Alexa recording their voice at times when they did not speak a wake word. However, they lack understanding that their voice recordings can be stored, accessed, and deleted. Although information about how to review and delete voice recordings is available to the users through their Alexa mobile application, efforts need to be made to educate the users about it.

REFERENCES

- [1] 2019. Manage Alexa Content and Devices. <https://www.amazon.com/mycd>.
- [2] Marwan Almaymoni Abdulaziz Alhadlaq, Jun Tang and Aleksandra Korolova. 2017. Privacy in the Amazon Alexa Skills Ecosystem. In *PETS*.
- [3] Amazon. 2019. Alexa and Alexa Device Terms, Warranties, and Notices. <https://www.amazon.com/gp/help/customer/display.html?nodeId=201602230>.
- [4] Amazon. 2019. Alexa Feature Help. <https://www.amazon.com/gp/help/customer/display.html?nodeId=201602040/>.
- [5] Malarie Gokey. 2018. The best smart speaker you can buy. <https://www.businessinsider.com/best-smart-speaker-amazon-echo>.
- [6] Stacey Gray. 2016. Always On: Privacy Implications of Microphone-Enabled Devices. In *Future of privacy forum*. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5714311/>.
- [7] Jeffrey Voas Hyunji Chung, Michaela Iorga and Sangjin Lee. [n. d.]. Alexa, Can I Trust You? <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5714311/>.
- [8] Carlos Jensen, Colin Potts, and Christian Jensen. 2005. Privacy practices of Internet users: self-reports versus observed behavior. *International Journal of Human-Computer Studies* 63, 1-2 (2005), 203–227.
- [9] Ruogu Kang, Stephanie Brown, Laura Dabbish, and Sara Kiesler. 2014. Privacy attitudes of mechanical turk workers and the us public. In *10th Symposium On Usable Privacy and Security (SOUPS) 2014*. 37–49.
- [10] Patrick Gage Kelley. 2010. Conducting usable privacy & security studies with amazon's mechanical turk. In *Symposium on Usable Privacy and Security (SOUPS) (Redmond, WA)*.
- [11] Bret Kinsella. 2017. Amazon Alexa Skill Count Passes 15,000 in the U.S. <https://voicebot.ai/2017/07/02/amazon-alexa-skill-count-passes-15000-in-the-u-s/>.
- [12] Ponnurangam Kumaraguru and Lorrie Faith Cranor. 2005. *Privacy Indexes: A Survey of Westin's Studies*. Technical Report CMU-ISRI-5-138. Institute for Software Research International, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA.
- [13] H. Lee and A. Kobsa. 2016. Understanding user privacy in Internet of Things environments. In *2016 IEEE 3rd World Forum on Internet of Things (WF-IoT)*. 407–412. <https://doi.org/10.1109/WF-IoT.2016.7845392>.
- [14] Yabing Liu, Krishna P. Gummadu, Balachander Krishnamurthy, and Alan Mislove. 2011. Analyzing Facebook Privacy Settings: User Expectations vs. Reality. In *Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference (IMC '11)*. ACM, New York, NY, USA, 61–70. <https://doi.org/10.1145/2068816.2068823>.
- [15] Emily McReynolds, Sarah Hubbard, Timothy Lau, Aditya Saraf, Maya Cakmak, and Franziska Roesner. 2017. Toys That Listen: A Study of Parents, Children, and Internet-Connected Toys. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, USA, 5197–5207. <https://doi.org/10.1145/3025453.3025735>.
- [16] David Nield. 2018. Alexa vs Siri vs Google Assistant vs Cortana vs Bixby. <https://www.popsi.com/digital-assistant-showdown>.
- [17] Anne Pfeifle. 2012. Alexa, What Should We Do About Privacy? Protecting Privacy for Users of Voice-Activated Devices. In *Washington Law Review*. <https://digital.law.washington.edu/dspace-law/bitstream/handle/1773.1/1778/93WLR0421.pdf>.
- [18] Franziska Roesner, Brian T Gill, and Tadayoshi Kohno. 2014. Sex, lies, or kittens? investigating the use of snapchat's self-destructing messages. In *International Conference on Financial Cryptography and Data Security*. Springer, 64–76.
- [19] Yang Wang, Gregory Norice, and Lorrie Faith Cranor. 2011. Who is concerned about what? A study of American, Chinese and Indian users' privacy concerns on social network sites. In *International Conference on Trust and Trustworthy Computing*. Springer, 146–153.
- [20] Matt Weinberger. 2017. Why Amazon's Echo is totally dominating. <https://www.businessinsider.com/amazon-echo-google-home-microsoft-cortana-apple-siri-2017-1>.
- [21] Allison Woodruff, Vasyli Pihur, Sunny Consolvo, Laura Brandimarte, and Alessandro Acquisti. 2014. Would a Privacy Fundamentalist Sell Their DNA for \$1000...If Nothing Bad Happened as a Result? The Westin Categories, Behavioral Intentions, and Consequences. In *10th Symposium On Usable Privacy and Security (SOUPS 2014)*. USENIX Association, Menlo Park, CA, 1–18. <https://www.usenix.org/conference/soups2014/proceedings/presentation/woodruff>.

A APPENDIX

A.1 Screening Question

- (1) Which of the following virtual assistant(s) do you own/use?
 - ☐ Amazon Echo, Echo Dot, Echo Show, Echo Spot, Echo Tap, Echo Plus
 - ☐ Google Now
 - ☐ Google Home
 - ☐ Siri
 - ☐ Cortana
 - ☐ Other
 - ☐ None

A.2 Demographic and Virtual Assistant Usage Questions

- (1) Which age group do you belong to?
 - ☐ 18 - 20
 - ☐ 20 - 30
 - ☐ 30 - 40
 - ☐ 40 - 50
 - ☐ 50 - 60
 - ☐ 60 and above
 - ☐ Decline to answer
- (2) What is your gender?
 - ☐ Male
 - ☐ Female
 - ☐ Other
 - ☐ Decline to answer
- (3) Which of the following best describes your highest achieved education level?
 - ☐ High school
 - ☐ 2 years of college/ Associate Degree
 - ☐ 4 years of college/ Undergraduate Degree
 - ☐ More than 4 years of college/ Graduate Degree
 - ☐ Decline to answer
- (4) Do you have a college degree or work experience in computer science, software development, web development, or similar computer-related fields?
 - ☐ Yes
 - ☐ No
 - ☐ Decline to answer
- (5) What is your ethnicity?
 - ☐ Asian
 - ☐ American Indian or Alaska Native
 - ☐ Hispanic, Latino, or Spanish origin
 - ☐ White
 - ☐ Black, African American

- ☐ Native Hawaiian or Other Pacific Islander
- ☐ Some other race or origin
- ☐ Decline to answer

(6) Which of the following Alexa-enabled device(s) do you own?

- ☐ Amazon Echo
- ☐ Amazon Echo Dot
- ☐ Amazon Echo Plus
- ☐ Amazon Echo Show
- ☐ Amazon Echo Spot
- ☐ Amazon Tap

(7) How many apps/skills have you linked to your virtual assistant?

(8) Please specify the duration of use of your Alexa-enabled device(s)

- ☐ 0 - 6 months
- ☐ 6 months - 1 year
- ☐ 1 year - 2 years
- ☐ 2 years - 3 years
- ☐ More than 3 years

(9) What wake word do you use for your Alexa-enabled device?

(10) The following question is for verification purposes to determine that you own an Alexa-enabled device. You will only be paid if you pass this validation check.

Please provide the 9 digit Device Software Version of your Alexa-enabled device by following these steps:

- (a) Open the application.
 - (b) Select Alexa Devices from the left menu.
 - (c) From the device list, select you Echo, Echo Dot, Echo Show, Echo Spot, Echo Tap or Echo Plus device
 - (d) Scroll down to About and find the Device Software Version
- _____

(11) Please state briefly what tasks you normally use your Amazon Alexa voice assistant for? e.g., turning on/off lights.

A.3 Alexa's Data Recording Storage, Access, and Deletion Behavior Related Questions

(1) In your opinion, does Amazon Alexa store any data? If yes, what data?

(2) In your opinion, where does Amazon Alexa store the data you specified?

(3) In your opinion, which persons/entities can potentially access this data?

(4) In your opinion, can you delete the data stored by Amazon Alexa?

- ☐ Yes
- ☐ No
- ☐ I am not sure

(5) Studies have shown that people do not pay attention to all questions in the surveys. This impacts the quality of the collected data. In order to ensure that you are paying attention, please select "Somewhat disagree" from the options below

- ☐ Strongly Agree
- ☐ Somewhat Agree
- ☐ Neither agree nor disagree
- ☐ Somewhat Disagree
- ☐ Strongly Disagree

A.4 Alexa's Unintended Voice Recording Awareness/Perceptions Questions

(1) Open the Alexa application inside your smartphone. Select settings from the bottom menu and then select History. Select Filter by date and choose May 1st as the start date and Aug 1st as the end date. From the instances you see in the filtered list, count the instances displaying "Text not available. Click to play recording". Approximately how many such instances appeared in your app? Please write a numerical value

(2) In order to verify that instances displaying "Text not available. Click to play recording" exist in your Alexa app, click on ONLY ONE such instance and take a screenshot of the window displaying it.

Note: You only need to take a screenshot of one of these instances. This window does not contain any personal information, and we only need its screenshot to verify that instances displaying "Text not available. Click to play recording" exist in your Alexa app.

Listen to a few of the instances that display "Text not available. Click to play recording" and answer the following questions.

(3) Do you remember giving voice commands that would result in these recordings?

- ☐ Yes, these recordings contain my commands and were intentional
- ☐ No, these recordings do not contain my commands and were unintentional

(4) Use the slider to express your feelings after listening to these recordings.

- ☐ Shocked
- ☐ Neutral
- ☐ Surprised

(5) What kind of information do these voice recordings contain?

(6) Do these voice recordings contain private/sensitive information?

- ☐ All of them
- ☐ Some of them
- ☐ None of them

(7) Do you consider these voice recordings a breach of privacy? In other words, the fact that voice was being recorded without your knowledge is an invasion of privacy.

- ☐ Strongly Agree
- ☐ Somewhat Agree

- ☐ Neither agree nor disagree
☐ Somewhat Disagree
☐ Strongly Disagree
- (8) Do you believe you can delete these voice recordings?
- ☐ Yes
☐ No
☐ I am not sure
- (9) If you selected "Yes" for the previous question, briefly explain how you would delete these recordings, otherwise type "N/A".
-
- (10) Studies have shown that people do not pay attention to all questions in the surveys. This impacts the quality of the collected data. In order to ensure that you are paying attention, please select 2 from the options below.
- ☐ 1
☐ 2
☐ 3
☐ 4
☐ 5
- (11) Please express how much you agree or disagree to the following statements.
- Consumers have lost all control over how personal information is collected and used by companies.
- ☐ Strongly Agree
☐ Somewhat Agree
☐ Neither agree nor disagree
- ☐ Somewhat Disagree
☐ Strongly Disagree
- Most businesses handle the personal information they collect about consumers in a proper and confidential way.
- ☐ Strongly Agree
☐ Somewhat Agree
☐ Neither agree nor disagree
☐ Somewhat Disagree
☐ Strongly Disagree
- Existing laws and organizational practices provide a reasonable level of protection for consumer privacy.
- ☐ Strongly Agree
☐ Somewhat Agree
☐ Neither agree nor disagree
☐ Somewhat Disagree
☐ Strongly Disagree
- (12) Please indicate how concerned you are about data privacy on a scale from 1 to 10.
- (Unconcerned) (Strongly concerned)*
- ☐ 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5 ☐ 6 ☐ 7 ☐ 8 ☐ 9 ☐ 10
- (13) If there is there anything else you would like to mention, please write it below.
-
- (14) Create an 8 character code consisting of letters and digits (e.g., 1qbn4hbl) and write it down here. Use this same code in the code box of your HIT submission window. DO NOT use 1qbn4hbl.
-