“Alexa, Stop Recording”: Mismatches between Smart Speaker Privacy Controls and User Needs

Josephine Lau
School of Information
University of Michigan
Ann Arbor, MI, USA
jlauuom@umich.edu

Benjamin Zimmerman
School of Information
University of Michigan
Ann Arbor, MI, USA
benzim@umich.edu

Florian Schaub
School of Information
University of Michigan
Ann Arbor, MI, USA
fschaub@umich.edu

Abstract
Smart speakers, like Amazon Echo and Google Home, provide benefits and convenience through their integrated voice assistants, but also raise privacy concerns due to their continuously listening microphones. We studied users' privacy-seeking behaviors around these devices and their use of current privacy controls. Through a diary study and in-home interviews with seventeen smart speaker users, we found users rarely engaged in privacy-seeking behaviors or utilized current privacy controls, which are currently not addressing their needs. Our findings can inform the design of privacy controls in future smart speakers.

Author Keywords
Privacy; smart speaker; privacy controls.

Introduction
The Amazon Echo smart speaker debuted in November 2014 [21]. Since then, Google [5, 8], Apple [2], and Microsoft [20] have introduced their own smart speakers. Millions of smart speakers have been sold [12] and worldwide spending on these devices is expected to reach $2 billion by 2020 [10].

Smart speakers offer users hands-free voice control, but to detect and respond to voice commands, the speakers' microphones must continuously be on to listen for their “*wake*

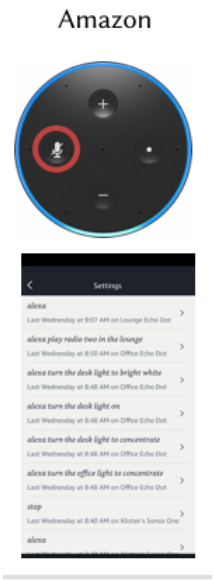

Figure 1: Privacy Controls for Amazon Echo and Google Home: mute buttons on top, audio logs on bottom.

word" [1, 11]. Considering that microphones are perceived as one of the most privacy-invasive sensors [7, 16] and homes are one of the most privacy-sensitive locations for Internet of Things (IoT) data collection [15], the privacy implications of smart speakers' 'always listening' capabilities have been the focus of much public debate [9, 4]

Prior studies have revealed how the detailed data collection from a single IoT device in the home [13] and the aggregation of sensed data [7] can reveal intimate insights on residents' activities. For example, a smart meter's power measurements can reveal when nobody is home [13]. Zeng et al. [22] included smart speakers in their study on security and privacy concerns in smart homes, but did not focus on smart speakers. They found that participants lacked security and privacy concerns regarding smart homes because they did not feel personally targeted, the trusted potential adversarial actors (like companies or governments), or they believed their existing mitigation strategies to be sufficient [22].

To address privacy concerns with audio/video recording technologies, researchers have recommended recording indicators [14] and user interfaces that display privacy risks and provide settings to control use and dissemination of collected data [7]. Such indicators and interfaces have to be designed carefully in order to be effective [17, 19, 18].

Current smart speakers are equipped with some privacy features. Speech recognition is performed locally by the device until the activation keyword has been detected, at which point the subsequent voice command is forwarded to the maker's servers for processing. In addition, most smart speakers are equipped with a physical button to mute the microphones. Companion mobile apps and websites enable users to review and delete prior voice interactions with the device (see Figure 1). However, consumers' use of these privacy controls and their privacy-seeking behaviors around the devices have not been studied in detail. An assessment of the usability of these controls and an understanding of current users' privacy-seeking behaviors could lead to improved privacy control designs.

Study Design

We conducted a diary study followed by semi-structured interviews with users in their homes to gain insights on their day-to-day behaviors around smart speakers, and to discern their privacy perceptions and concerns regarding smart speakers. Our study was approved by our institution's IRB.

Following the diary-interview method [23, 3], we first conducted a one-week diary study with smart speaker users. We chose this method to learn how users engage with their smart speakers on a daily basis and to mitigate recall bias often present in isolated interviews [23]. Users were asked to submit at least one diary entry per day for one week through an online survey. We asked about instances in which they used the device, as well as times they had considered using their smart speaker but did not – those situations might signify privacy-seeking behaviors. We also asked users to report accidental smart speaker activations, in case these instances triggered reflection on privacy. Users could submit entries for each instance or daily summaries; they could also indicate that they had no speaker interaction on a particular day. Once participants started the diary study, we tracked their progress throughout the week via the survey tool and sent email reminders if they had not submitted at least one entry by 9pm. Participants were compensated $15 for completing the diary study.

After completing the diary study, we interviewed each user

in their home. We modeled parts of our interview script after other qualitative smart home and IoT privacy studies [14, 22, 6]. A user's diary entries served as prompts for interview questions, especially entries about non-interaction or accidental interaction. If privacy did not come up naturally, we explicitly asked about users' privacy perceptions and concerns regarding smart speakers. We probed to assess their awareness and use of current speakers' privacy controls including the physical mute button and audio logs. We further asked whether users considered smart speakers' current privacy controls sufficient, and how they would want a 'dream speaker' to protect their privacy. Participants received another $15 for completing the interview.

Recruitment and Demographics

We recruited participants through flyers, announcements on Facebook and the local subreddit, and an email to a university mailing list. Over fifty potential users completed our online screening survey. Based on screening survey responses, we invited users who reported use of at least one smart speaker. Seventeen users completed both portions of the study. The diary study was conducted in July 2017 and interviews were conducted shortly after. Interviews with users lasted 25–59 minutes (Median: 44 minutes). Users who had owned their speaker for less than a month (3) were balanced out by those who had owned theirs for over a year (4); other users were evenly distributed in between.

Findings

Privacy Perceptions of and Concerns with Smart Speakers

To contextualize our design suggestions, we briefly discuss users' privacy perceptions and concerns with smart speakers. When probed about privacy, almost all users said they were not concerned, with some being more confident in this assertion than others; two users never considered privacy an issue about which they should be concerned. Many justified their lack of privacy concern by labeling themselves as "not interesting" or having "nothing to hide." They also rationalized that it would require 'too much storage' and 'too much processing power' to collect, store, and analyze 'everything' from their smart speakers. However, these participants' assumptions miss the possibility that companies could feasibly collect everything about specific individuals when compelled to do so. Two users considered the information they give to their smart speakers to be small additions to the body of knowledge companies already have about them. This perception illustrates a resignation to the loss of control over privacy.

Lack of Privacy-Seeking Behaviors

Since prior research had found that microphones were perceived as one of the most intrusive sensors in home contexts [7, 16], we studied if and how users were engaging in privacy-seeking behaviors around smart speakers. However, none of our users reported any privacy-seeking behaviors in their diaries. In interviews, we probed deeper to see whether they might switch to another device (e.g. their phone) for more sensitive interactions, or whether accidental smart speaker activations might make users worry about being listened to. We found that rarely any users engaged in privacy-seeking behaviors around their speakers, and the six users who reported accidental activations of their speakers found these interactions amusing and not privacy-concerning.

Non-use of Current Privacy Controls

Many users were aware of the ability to review audio logs and press the mute button on their smart speaker, but for multiple reasons, they did not use those controls for privacy regulation. Four users expressed that using the mute button would negate the device's primary functionality – hands-free use. Additionally, despite being aware of the mute

button, a few participants incorrectly believed they could silence their speaker through a voice command, such as *"Alexa, mute yourself"* (U13) or *"Hey Google, stop recording"* (U11). And although most users were aware of audio logs, only one deleted a log for privacy reasons. We noted that while users were aware of the ability to look back on their audio logs in the companion apps, very few made the connection that this capability could be used as form of privacy control. Indeed, many of our users said it was not common for them to look at the phone app or the audio logs – they saw no reason to. Two primary users used their audio logs, not as privacy controls, but rather to monitor secondary and incidental users.

Unmet Privacy Control Needs

The non-use of smart speakers' privacy controls suggests that users might not feel the need to regulate privacy around those devices. That is not the case. Multiple users expressed privacy control needs that were just not met by existing controls. For example, several users mentioned using private browsing modes while searching or browsing the Internet, and two users pointed out how there is no similar functionality on smart speakers. This shows that users desired a more proactive way to prevent their recordings from being saved; deleting an audio log is retroactive and requires users to remember to go back and look at the logs. Users also talked about privacy issues caused when their smart speakers were utilized by multiple users. One user talked about how their Echo's limit of two adults to a household has forced her and her roommate to share credentials, which may lead to privacy issues in the future. Additionally, she was concerned about the speaker leaking sensitive information, but didn't see how she could mitigate that.

Conclusion

While smart speakers can offer users convenience and more efficient lifestyles, their use comes with privacy risks and implications. In our study, we found that not all users were aware of these risks or underestimated the risks involved in using this technology. While most were aware of current smart speakers' two privacy controls, they were not actively using them. From this, we propose that designers integrate privacy controls into speaker's conversational capability, facilitate proactive rather than retroactive control, and set privacy-friendly defaults.

Acknowledgments

This research was partially supported by the Institute of Museum and Library Services through grant no. LG-06-14-0122-14 ("Research Experience for Masters Students") and by the University of Michigan School of Information.

REFERENCES

1. Amazon. 2018. Amazon Echo (2nd Generation) — Always Ready, Connected, and Fast. Just Ask. (2018). https://www.amazon.com/Generation-improved-sound-powered-design/dp/B06XCM9LJ4 Accessed Jan 7, 2018.
2. Apple. 2017. HomePod Reinvents Music in the Home. (May 2017). https://www.apple.com/newsroom/2017/06/homepod-reinvents-music-in-the-home/
3. Ruth Bartlett and Christine Milligan. 2015. *What Is Diary Method?* Bloomsbury Publishing.
4. BBC. 2017. Amazon Hands over Echo 'murder' Data. *BBC News* (March 2017). http://www.bbc.com/news/technology-39191056
5. Dieter Bohn. 2016. Google Home: A Speaker to Finally Take on the Amazon Echo. *The Verge* (May 2016).

https://www.theverge.com/2016/5/18/11688376/google-home-speaker-announced-virtual-assistant-io-2016

6. A.J. Bernheim Brush, Bongshin Lee, Ratul Mahajan, Sharad Agarwal, Stefan Saroiu, and Colin Dixon. 2011. Home Automation in the Wild: Challenges and Opportunities. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. ACM, New York, NY, USA, 2115–2124. DOI: http://dx.doi.org/10.1145/1978942.1979249

7. J. Bugeja, A. Jacobsson, and P. Davidsson. 2016. On Privacy and Security Challenges in Smart Connected Homes. In *2016 European Intelligence and Security Informatics Conference (EISIC)*. 172–175. DOI: http://dx.doi.org/10.1109/EISIC.2016.044

8. Rishi Chandra. 2017. Welcoming Mini and Max to the Google Home Family. (Oct. 2017). https://www.blog.google/products/home/welcoming-mini-and-max-google-home-family/

9. Adam Clark Estes. 2017. Don't Buy Anyone an Echo. (May 2017). https://gizmodo.com/dont-buy-anyone-an-echo-1820981732

10. Gartner. 2016. Gartner Says Worldwide Spending on VPA-Enabled Wireless Speakers Will Top $2 Billion by 2020. (Oct. 2016). http://www.gartner.com/newsroom/id/3464317

11. Google. 2018. Google Home. (2018). https://store.google.com/product/google_home Accessed Jan 7, 2018.

12. Jacob Kastrenakes. 2018. Google Sold over 6 Million Home Speakers since Mid-October. (Jan. 2018). https://www.theverge.com/2018/1/5/16855982/google-home-sales-figures-holidays-2017

13. Andreas Kirmse. 2012. *Privacy in Smart Homes*. Technical Report. http://kirmandi.rumeln.net/data/paper-Privacy.in.Smart.Homes.pdf

14. 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. DOI: http://dx.doi.org/10.1145/3025453.3025735

15. Pardis Emami Naeini, Sruti Bhagavatula, Hana Habib, Martin Degeling, Lujo Bauer, Lorrie Cranor, and Norman Sadeh. 2017. Privacy Expectations and Preferences in an IoT World. In *Proceedings of the Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)*. Santa Clara, CA. https://www.usenix.org/system/files/conference/soups2017/soups2017-naeini.pdf

16. Antti Oulasvirta, Aurora Pihlajamaa, Jukka Perkiö, Debarshi Ray, Taneli Vähäkangas, Tero Hasu, Niklas Vainio, and Petri Myllymäki. 2012. Long-Term Effects of Ubiquitous Surveillance in the Home. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing (UbiComp '12)*. ACM, New York, NY, USA, 41–50. DOI: http://dx.doi.org/10.1145/2370216.2370224

17. Rebecca S. Portnoff, Linda N. Lee, Serge Egelman, Pratyush Mishra, Derek Leung, and David Wagner. 2015. Somebody's Watching Me?: Assessing the Effectiveness of Webcam Indicator Lights. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 1649–1658. DOI: http://dx.doi.org/10.1145/2702123.2702164

18. F. Schaub, R. Balebako, and L. F. Cranor. 2017. Designing Effective Privacy Notices and Controls. *IEEE Internet Computing* 21, 3 (May 2017), 70–77. DOI: http://dx.doi.org/10.1109/MIC.2017.75

19. Florian Schaub, Rebecca Balebako, Adam L. Durity, and Lorrie Cranor. 2015. A Design Space for Effective Privacy Notices. In *Eleventh Symposium On Usable Privacy and Security*. Ottawa, 1–17. https://www.usenix.org/conference/soups2015/proceedings/presentation/schaub

20. Andrew Shuman. 2017. Hey Cortana, Set a Reminder: Harman Kardon Invoke Voice-Activated Speaker Available October 22. (Oct. 2017). https://blogs.windows.com/windowsexperience/2017/10/20/harman-kardon-invoke-voice-activated-speaker-available-october-22/

21. Chris Welch. 2014. Amazon Just Surprised Everyone with a Crazy Speaker That Talks to You. (Nov. 2014). https://www.theverge.com/2014/11/6/7167793/amazon-echo-speaker-announced

22. Eric Zeng, Shrirang Mare, and Franziska Roesner. 2017. End User Security and Privacy Concerns with Smart Homes. In *Proceedings of the Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017*. Santa Clara, CA.

23. Don H. Zimmerman and D. Lawrence Wieder. 1977. The Diary: "Diary-Interview Method". *Urban Life; Newbury Park, Calif.* 5, 4 (Jan. 1977), 479–498. http://search.proquest.com/docview/1292940175/citation/A16505C65BA645C4PQ/1