Use Case: Passively Listening Personal Assistants

Nathan Malkin, Primal Wijesekera, Serge Egelman, David Wagner
University of California, Berkeley
{nmalkin,primal,egelman,daw}@cs.berkeley.edu

As the Internet of Things grows, it will continue to provide both conveniences and significant privacy challenges. One existing category of devices that exemplifies these tradeoffs is personal voice assistants. These devices—such as Apple's Siri, Google's Home, and Amazon's Alexa—respond to users' commands, such as sending a message, performing a query, or activating another smart-home device. In their most popular form factor, they work as a smart speaker, always listening for "trigger words" (such as "Hey Siri" or "Ok Google"), then recording and analyzing any audio that follows it, in order to extract (and act on) the user's instructions.

Even the current setup, with users explicitly triggering the device, presents privacy problems. Most people don't understand when a smart device is listening and where it is sending data [2]. Some may not even realize that their recordings are stored in the cloud—though law enforcement certainly does, and has requested the collected data as part of investigations [3]. Advertisers have also sought to exploit the insights these devices offer into shoppers' lives. Patent filings from Amazon and Google have described designs for using data they collect from smart speakers for targeted advertising [1].

One notable assumption in the patent filings is that the devices are constantly listening and analyzing data. In fact, passively listening (or passively watching) devices are already on the market. As one example, the Google Clips camera takes pictures continuously and then selects the "best" photos among them [4]. Many other in-home cameras are designed and marketed for security purposes. As their capabilities expand, it is highly likely that more people will welcome them into their homes, despite potential privacy concerns.

But what does it mean for consumers' privacy if a device is always listening? These passive listening devices will not need a trigger word to start capturing the audio surrounding the device. With the trigger word gone, users have even less transparency into when these devices are capturing or processing audio. Understanding privacy expectations is critical in such a scenario, because the chances of the device capturing audio when it is not expected become significantly higher.

Lack of transparency won't be the only privacy issue for future passive listening devices. We envision the ecosystem surrounding these devices following the path of smartphones, where the role of the operating system will be as more of a mediator. We believe more third-party applications will be able to run on these devices providing narrowly defined functions such as calendaring, ride sharing, etc. Already, voice assistants support third-party applications (for example, "Skills" for Alexa). What will be the privacy protections when these devices shift to passive listening?

Will the third parties also get full access to audio from within our homes? With the reduced transparency, the possibility of third parties accessing sensitive data when it is not expected is not a desirable outcome, and even the platforms themselves might want to impose restrictions on the data they collect. The question—and challenge—is how. How can a passively listening device identify when it should or shouldn't be listening? How can a person specify which conversations are fair game for an app and which are private? Can a conversation be appropriate for one app while being the opposite for another?

Answering many of these questions is a job for contextual integrity, as less powerful definitions of privacy are insufficient for this situation. For example, there are no conversations that can simply be labeled as private. Instead, it is necessary to consider the information flows, people's expectations, and contextual norms.

While we know that contextual integrity is a useful framework for modeling people's expectations, there are also challenges in applying it to passively listening devices. In particular, what are the operative transmission principles? How can the "context" be identified, and which set of contextual norms apply to a given situation?

Passive listening devices also pose a different set of challenges to CI, compared to how the framework has been used for smartphones. In applying contextual integrity to the passive listening use case, we are likely to encounter questions that will require a deeper understanding of the surrounding context of a given conversation (i.e., data flow). For example, when a conversation is taking place, can it be in more than one context at the same time? Can different people involved in the same conversation have different expectations or perceive the same conversation with a different contextual perspective?

Answering these questions is a necessary step towards being able to apply the CI framework to managing privacy in this scenario, as well as in many related ones.

REFERENCES

[1] Consumer Watchdog. 2017. Home Assistant Adopter Beware: Google, Amazon Digital Assistant Patents Reveal Plans for Mass Snooping. http://www.consumerwatchdog.org/index.php/privacy-technology/home-assistant-adopter-beware-google-amazon-digital-assistant-patents-reveal

[2] Nathan Malkin, Julia Bernd, Maritza Johnson, and Serge Egelman. 2018. What Can't Data Be Used For? User Privacy Expectations about Smart TVs. In *Proceedings of the 3rd European Workshop on Usable Security*.

[3] Christopher Mele. 2016. Bid for Access to Amazon Echo Audio in Murder Case Raises Privacy Concerns. *The New York Times* (Dec. 2016). https://www.nytimes.com/2016/12/28/business/amazon-echo-murder-case-arkansas.html

[4] Ben Popper. 2017. Google's new Clips camera is invasive, creepy, and perfect for a parent like me. *The Verge* (Oct. 2017). https://www.theverge.com/2017/10/5/16428708/google-clips-camera-privacy-parents-children

Symposium on Applications of Contextual Integrity, 2018