
eGregor: An Eldritch Privacy Mental Model For Smart Assistants

Christine Geeng

University of Washington
Seattle, WA, USA
christinegeeng@gmail.com

Anonymous Author

Abstract

The increasing popularity of smart personal assistants has meant the rapid inclusion of data-collecting technology in homes. Research has shown that the privacy notices for these smart devices can be ineffective, as users often have incorrect mental models about what happens to data collected from them. To provide more effective data collection cues, we present a redesign of traditional, friendly smart assistant personas: eGregor, an eldritch hive mind being. Using science fiction concepts in conjunction with visceral notice, a concept that eschews purely text-based privacy indicators, eGregor more clearly represents the data practices of a hypothetical parent company. It has various aesthetic and auditory indicators and an intuitive and terrifying persona. We draw attention to the potential for major smart personal assistant companies to improve upon their user interface designs.

Author Keywords

Smart home; smart assistant; privacy; notice and consent; ubiquitous computing.

Introduction

Smart personal assistants (SPAs) such as Alexa, Google Home, and Siri have become popular in domestic life; over 39 million adults in the U.S. now own a smart assistant speaker [3]. SPAs provide significant value to consumers

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 the author(s) 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.

Copyright held by the owner/author(s). Publication rights licensed to ACM.
CHI'20 Extended Abstracts, April 25–30, 2020, Honolulu, HI, USA
ACM 978-1-4503-6819-3/20/04.
<https://doi.org/10.1145/3334480.3381827>

with hands-free control, Internet search, and other functionality. Due to both their popularity and their consumption of user data, these devices have attracted the attention of privacy and security researchers.

While researchers may understand the privacy and security implications of owning an SPA, users often do not. Studies both by Abdi et al. and Zeng et al. found that user mental models of smart assistants are often incomplete or wrong entirely [1, 21]. These researchers have identified a need to better communicate data usage and collection practices of assistants through design.

Inspired by the concept of “visceral notice” to better indicate data collection practices to users [6], we propose making over the smart assistant aesthetic using cultural references to help consumers conceptualize its capabilities. We propose eGregor, a hive mind Eldritch horror voice assistant persona analogous to Siri, Alexa, or Cortana, based on the occult concept of egregore as a collective consciousness. We identify several unexpected smart assistant practices and describe how our system redesign provides clearer, more visceral notices for these use cases.

SPA Privacy and Poor Mental Models

SPA Architecture

A typical SPA architecture consists of an Internet-connected device that may process some commands on the device itself, and others commands (such as requests for real-time information) on an SPA cloud service (see Figure 2). Vocal command recordings are stored both locally on device and on the SPA company servers [9].

Third-party skills, like calling an Uber or turning off a Philips Hue light bulb, involve sending commands to third-party servers for processing. These external companies may have different data policies than SPA providers do. Abdi et

al. noted that participants did not know how long their SPA stored their data, and some were unaware that the data ever left the device. No participants had considered that their data could be stored on third-party servers as well [1].

Previously Amazon, Google and Apple have “hire[d] contractors to listen to anonymized user audio clips for the purposes of improving their respective assistant’s capabilities” [7]. Google and Apple have since suspended this practice, and Amazon now allows users to opt out of human review [5] – although opt-out requires navigation to the website’s settings page, which is decoupled from the device itself.

Notice and Consent

The dominant corporate paradigm for data collection practices is “notice and consent”: SPA users read a privacy policy notice that “is a public document from an organization that explains how that organization processes personal data and how it applies data protection principles” [20], and then check a box to consent to this policy. Companies which rely on this posit that users have sufficient information to give consent for data collection; however, consumers almost never read privacy policies [14]. Moreover, a consumer who does read a privacy policy is likely to find confusing legal jargon.

Privacy notices also exist outside of policy text on a website. For example, it is common for an LED to glow while a camera is recording video. According to Schaub et al., notices are often ineffective because they are overly complex, don’t offer real choices, fatigue users, or are decoupled from the actual device [16]. The last issue is particularly troublesome for smart home devices, as sensors in ubiquitous computing environments are often disassociated from a medium to indicate data recording (such as a screen) [4, 22]. This can be an issue not only for smart device owners,

but also for bystanders, or for incidental users, who may not even know that a device is nearby [11].

Calo suggested “visceral notice” as an improvement upon purely language-based privacy notice and its issues relayed above; he suggests other ways to experience warnings, such as cellphones using shutter noises when taking photos despite not having a physical shutter [6].

Given that consumers may have incorrect mental models of what happens to their data with smart home devices, particularly incidental users who did not purchase devices but are active near them [21, 10], we designed a new SPA persona to make notice on data capture, storage, access, and use extremely visceral. Our guiding principles in this redesign are as follows:

- You should know when you are being watched.
- You should know when your data is being shared.
- You should know that your data may exist forever.
- Your SPA is not your friend; it is a corporate entity. It is alien.

In the remainder of this paper, we introduce our creation and discuss its cosmic superiority to existing models.

eGregor: A Privacy-Notice-Conscious SPA

eGregor is a voice assistant with the persona of an elder being from some unfathomable dimension or decaying kingdom. It provides all the functionality that Siri, Alexa, and Google Home typically provide. eGregor can set timers, control your smart lights, search the web, and check your calendar. It is equipped with a voice that sounds like a million supplicants bellowing in agonized unison. It can make



Figure 1: eGregor is a cosmic horror smart personal assistant that has a microphone, speaker, and several eyes.

phone calls, and it can read texts, and it is covered in eyes. We believe that eGregor encourages a mental model consistent with actual privacy practices and risks endemic to SPA companies. Furthermore, eGregor improves upon the notice-and-consent paradigm through aesthetic, persistent, and verbal affordances indicating privacy practices on both a conscious and a visceral level.

Below we provide concrete examples of how eGregor improves upon standard data collection indicators.

Aesthetic Indicators

It is hard not to like your Alexa. Her informal style, unobtrusive design, and sheer utility rival that of a human personal assistant. She can remind you to wish your boss a happy birthday. She can read you directions on a road trip. She can even tell a joke or two if you ask nicely. There is

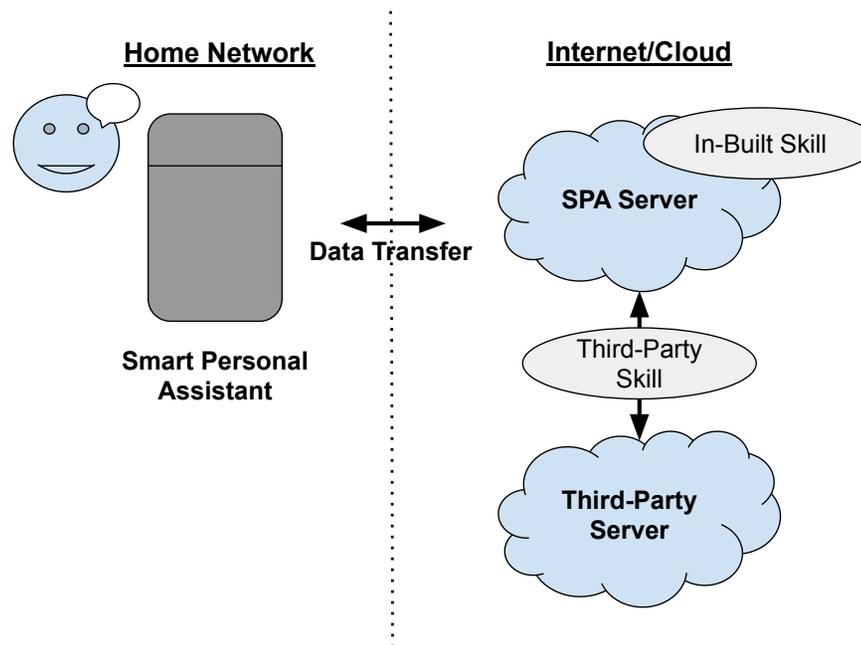


Figure 2: Diagram of SPA architecture and data flow (based on [1]).

Current Smart Assistants	eGregor
Hey Siri. OK Google. Hey Alexa.	eGregor, hear me! I relinquish my very voice to you oh Eternal Collector!
Alexa, delete everything I said today.*	eGregor, purge myself from your eyes, and I pray you will forget my puny existence!
Hey Siri, what's it like outside? OK Google, what's the chance of rain today? Alexa, what's the weather?	O, Great and Terrible Aggregator Of Timeless Knowledge! I consent to your discovery of my physical form if only you tell me what the coming tempest shall bring.

Table 1: Example dialogue between the user and eGregor. *Alexa may still retain transcript data even after this request [13].

something pleasantly person-like about Alexa, but she is not a human being. No matter how secure her protocols and benevolent her creator's intentions, Alexa does not care about you. She is a massively complicated, non-local, unfeeling marvel of statistics and software engineering. You do not understand her.

eGregor does not cloak the alien-ness of its design in the servile facade of a human personality. It celebrates its own inhumanity. eGregor presents itself not as a human individual, but as a vast and unknowable alien collective. This persona is meant to encourage a healthy mental model of the assistant and its use of your data in two ways. First, eGregor does not and cannot care about you. Second, you do not understand on a case-by-case basis what eGregor does with your data.

Stylistic elements borrowed from the genre of cosmic horror help to establish eGregor's alien indifference to the user. The device itself is organic in form and is plastered with an unsettling number of eyes, none of which form a matching pair. The device encourages users to treat it as superior to them. Furthermore, sound cues during recording are uncanny or occult by design.

eGregor demonstrates its storage of user data by presenting itself as a hive mind. It stores data locally on the device, on company servers, and possibly third-party servers where it may potentially get sold to data brokers. This spread of data is mirrored by eGregor's multiple consciousnesses. When the user asks eGregor to delete their audio recordings, while the audio is scrubbed from its physical form, the transcripts may remain saved in its other consciousnesses (see Table 1).

Even if eGregor were never to share your data outside of itself, such as with third-parties, it is not truly kept private if

eGregor is composed of uncountable individuals. The hive mind mental model mirrors the structure of real voice assistant companies, which are composed of many human individuals. Some individuals may listen to user voice recordings (with location data removed [8]) to improve eGregor's voice recognition if the user has not opted out of this practice [7]. The hive mind component of eGregor's presentation is relayed by its voice, which is composed of many distinct voices speaking unison, and its usage of the pronouns *we/us/ours* when speaking.

Persistent Recording Indicators

Users who are recorded without their awareness risk sharing private data by accident; thus any user interacting with eGregor should understand when their data is being captured. Of course, users will not understand eGregor itself, because eGregor is incomprehensible to our paltry human minds. Alexa, Google Home, and Siri emit polite audio chimes at the beginning and end of each recording period, and activate either an on-screen overlay (Siri) or a subtle LED bank (Google Home, Alexa) during recording. These signals, however, are not designed for maximum visibility or audibility to bystanders.

eGregor is not subtle. Its design includes audio and visual indications of recording which are persistent through the entire recording period. Visually, the many eyes dotting eGregor's surface blink open and closed during while the device records. The symbolic usage of eyes as privacy reminders has previously been shown to help users limit their data exposure [17]. To protect users who may not have a direct line-of-sight to the device, eGregor also emits diffuse red light which pulses at a regular rate. Consequently, while the device is recording, a vibrant red glow will illuminate nearby objects, increasing the chance that a bystander will notice the device and investigate.

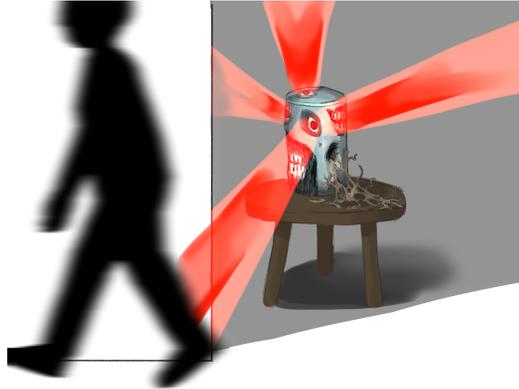


Figure 3: eGregor emits pulsing red light to provide indication of its exalted existence and data collection to bystanders who may not know a smart device is in the vicinity.

In addition to these visual cues, eGregor emits a loud audible signal during its entire recording period. The signal is composed of two overlaid audio tracks. The base layer is the so-called "Shepard Tone," which presents the illusion of a constantly rising pitch that never increases in frequency [18]. Our intention in choosing this illusion was to remind the user that their data is constantly ascending into the cloud, and therefore away from their control. The second layer of eGregor's persistent audio notification of recording is the sound of an occult choir ominously chanting eGregor's privacy policy. In this way users are reminded both that there may be multiple people listening to their recordings and that they have already granted eGregor legal access to their data.

Verbal Indicators

eGregor offers an improved suite of verbal commands that help users internalize eGregor's information flows (see Table 1). Users can use the wake word "eGregor" or use one of its formal titles: "the Eternal Collector," "the Undying Aggregator," "the Unending Archivist of Petty Truths," or "the Timeless Devourer of Data." eGregor's titles reinforce the fact that user data may never be deleted, even upon request. In summoning eGregor by title, the user must indicate knowledge of the longevity of their data. In this way eGregor – in a truly eldritch display – appropriates the user's own body and voice to provide notice of its privacy practices.

Users of eGregor are also encouraged to provide per utterance consent for the processing of their data. Users who successfully specify what data eGregor can use are rewarded with a curt "it is done," or a cool "we accept" from the smart assistant upon the completion of their utterance. Users who do not successfully consent to data collection on a per-utterance basis will be informed of what data is being taken in a long-form response berating the user for their insufficiency. eGregor still processes the user's data in these cases of course, because data is delicious to eGregor, and eGregor is always voracious.

Discussion

One could argue that this kind of redesign for a smart assistant would deter people from buying it due to its creepiness or its potential for setting a mood of malintent. This may be true; but with eGregor, consumers get the benefits of better indications of data policy that can help them make more informed decisions on smart home products they buy.

The risk of friendly SPA personas is that they elicit trust without necessarily providing a full picture of what com-

panies do with user data. Even information collected for a necessary business purpose may be processed for an undisclosed reason or shared without the user's knowledge in the pursuit of profit and technology development. eGregor, for all its improvements to providing notice, also does not fully inform users of all its intents, but its discomfoting presentation does not encourage unwarranted trust. Recognizing that even eGregor, the Unfading Devourer of Forbidden Knowledge, cannot fully inform users begs the question: *how can stakeholders in smart assistants better protect consumers from potential harms?*

In the area of privacy and law there have been some improvements, as GDPR has provided the groundwork and incentive for companies to comply with law and give consumers more control over their data. (Although mistakes can still happen with data, such as when Amazon Alexa voice recordings which were requested under GDPR were sent to the wrong user [19].) In design, Karmann et al. have produced Project Alias, an add-on to SPAs that gives users better privacy controls and ensures that SPA microphones do not listen outside of wake word commands, showing it is possible for SPA design to be both usable and privacy enhancing [12].

One limitation of improved, more visceral privacy indicators for smart home devices is that it is an individualistic take on a solution to potential data misuse and exploitation, one that is reliant on the user to understand both company privacy policies and how a smart device's underlying architecture works. No matter how compelling the notice is, users must still relinquish control over their data for their device to function. Invoking discomfort in the user does not change the nature of their contract with a corporation; it only makes them aware of their relative impotence.

Legislation like GDPR is a step in the right direction, as

trust in law can substitute trust in smart devices. Aphorpe et al. found that parents were more accepting of data-collecting Internet-connected toys if they were compliant with the U.S. Children's Online Privacy Protection Act (COPPA), though the researchers suggested COPPA may provide a false sense of security for children's privacy as "COPPA guidelines are relatively broad" [2].

For all of our critique, this paper is not meant to discourage the creation of smart personal assistants. SPAs offer a lot of helpful functionality, and immense potential "to provide inclusive, accessible interaction for people with a range of disabilities" [15]. Audio interactions allow for more efficient multi-tasking for vision-impaired users. Siri and Alexa and Google can provide an Internet of knowledge with a simple "Hey [blank]" and a question. But if billion-dollar companies have the resources to engineer smart devices, then they have the resources to better design privacy indicators to alert consumers. The power and data these companies hold is vast, but what they do with it is not unknowable.

Conclusion

We present eGregor, an Eldritch horror alternative to standard friendly smart personal assistant personas. It provides more visceral notice for users regarding when their data is captured, where it is stored, how long it is stored, and who has access to it, using aesthetic and auditory indicators.

Acknowledgements

We are grateful to the reviewers for their valuable feedback on our work. We also thank William Agnew and Franziska Roesner for their support and Atti Waller for her artistic input. We were heavily inspired by the works of Wayne Douglas Barlowe.

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 *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*. USENIX Association, Santa Clara, CA. <https://www.usenix.org/conference/soups2019/presentation/abdi>
- [2] Noah Apthorpe, Sarah Varghese, and Nick Feamster. 2019. Evaluating the Contextual Integrity of Privacy Regulation: Parents' IoT Toy Privacy Norms Versus COPPA. In *28th USENIX Security Symposium (USENIX Security 19)*. USENIX Association, Santa Clara, CA, 123–140. <https://www.usenix.org/conference/usenixsecurity19/presentation/apthorpe>
- [3] Scott Bay. 2018. AI assistants are poised for major growth in 2018. (2018). <https://venturebeat.com/2018/01/22/ai-assistants-are-poised-for-major-growth-in-2018/>.
- [4] Victoria Bellotti and Abigail Sellen. 1993. Design for Privacy in Ubiquitous Computing Environments. *Proceedings of the Third European Conference on Computer-Supported Cooperative Work ECSCW '93* (1993), 77–92. DOI: <http://dx.doi.org/10.1017/CB09781107415324.004>
- [5] Dieter Bohn. 2019. Amazon will let you opt out of human review of Alexa recordings. (2019). <https://www.theverge.com/2019/8/2/20752418/amazon-alexa-human-review-recordings-opt-out-eu>.
- [6] Ryan Calo. 2013. Against Notice Skepticism in Privacy (and Elsewhere). *Notre Dame Law Review* 87, 3 (2013). <https://digitalcommons.law.uw.edu/faculty-articles/29>
- [7] Ry Crist. 2019. Amazon and Google are listening to your voice recordings. Here's what we know about that. (2019). <https://www.cnet.com/how-to/amazon-and-google-are-listening-to-your-voice-recordings-heres-what-we-know>.
- [8] Anthony Cuthbertson. 2019. Google Defends Listening To Private Conversations On Google Home – But What Intimate Moments Are Recorded? (2019). <https://www.independent.co.uk/life-style/gadgets-and-tech/news/google-home-recordings-listen-privacy-amazon-alexa-hack-a9002096.html>.
- [9] Jide S. Edu, Jose M. Such, and Guillermo Suarez-Tangil. 2019. Smart Home Personal Assistants: A Security and Privacy Review. (2019).
- [10] Jason Hong. 2016. Toward a Safe and Secure Internet of Things. June (2016). Retrieved 2018-08-31 from <https://www.newamerica.org/cybersecurity-initiative/policy-papers/toward-a-safe-and-secure-internet-of-things/>
- [11] Roberto Hoyle, Robert Templeman, Steven Armes, Denise Anthony, David Crandall, and Apu Kapadia. 2014. Privacy Behaviors of Lifeloggers Using Wearable Cameras. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '14)*. Association for Computing Machinery, New York, NY, USA, 571–582. DOI: <http://dx.doi.org/10.1145/2632048.2632079>
- [12] Bjørn Karmann. 2018. Project Alias. (2018). http://bjoernkarmann.dk/project_alias.

- [13] Makena Kelly and Nick Statt. 2019. Amazon confirms it holds on to Alexa data even if you delete audio files. (2019). <https://www.theverge.com/2019/7/3/20681423/amazon-alexa-echo-chris-coons-data-transcripts-recording-privacy>.
- [14] Jonathan A Obar and Anne Oeldorf-Hirsch. 2018. The biggest lie on the internet: Ignoring the privacy policies and terms of service policies of social networking services. *Information, Communication & Society* (2018), 1–20.
- [15] Alisha Pradhan, Kanika Mehta, and Leah Findlater. 2018. “Accessibility Came by Accident”: Use of Voice-Controlled Intelligent Personal Assistants by People with Disabilities. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. Association for Computing Machinery, New York, NY, USA, Article Paper 459, 13 pages. DOI: <http://dx.doi.org/10.1145/3173574.3174033>
- [16] Florian Schaub, Rebecca Balebako, Adam L. Durity, and Lorrie Faith Cranor. 2015. A Design Space for Effective Privacy Notices. In *Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*. USENIX Association, Ottawa, 1–17. <https://www.usenix.org/conference/soups2015/proceedings/presentation/schaub>
- [17] Roman Schlegel, Apu Kapadia, and Adam J Lee. 2011. Eyeing your exposure: quantifying and controlling information sharing for improved privacy. In *Proceedings of the Seventh Symposium on Usable Privacy and Security*. ACM, 14.
- [18] Roger N Shepard. 1964. Circularity in judgments of relative pitch. *The Journal of the Acoustical Society of America* 36, 12 (1964), 2346–2353.
- [19] Nick Statt. 2019. Amazon sent 1,700 Alexa voice recordings to the wrong user following data request. (2019). <https://www.theverge.com/2018/12/20/18150531/amazon-alexa-voice-recordings-wrong-user-gdpr-privacy-ai>.
- [20] Ben Wolford. 2019. Writing a GDPR-compliant privacy notice. (2019). <https://gdpr.eu/privacy-notice>.
- [21] Eric Zeng, Shirang Mare, and Franziska Roesner. 2017. End user security and privacy concerns with smart homes. In *Thirteenth Symposium on Usable Privacy and Security ({SOUPS} 2017)*. 65–80.
- [22] Serena Zheng, Marshini Chetty, and Nick Feamster. 2018. User Perceptions of Privacy in Smart Homes. (2018). <http://arxiv.org/abs/1802.08182>