Carnegie Mellon University

Enhancing Transparency in AI Voice Assistants

Sponsor: Courtney Yang, Meta
Advisor: Lorrie Cranor

Ruiyang Liu
Prahaladh Chandrahasan
Ashutosh Sahu

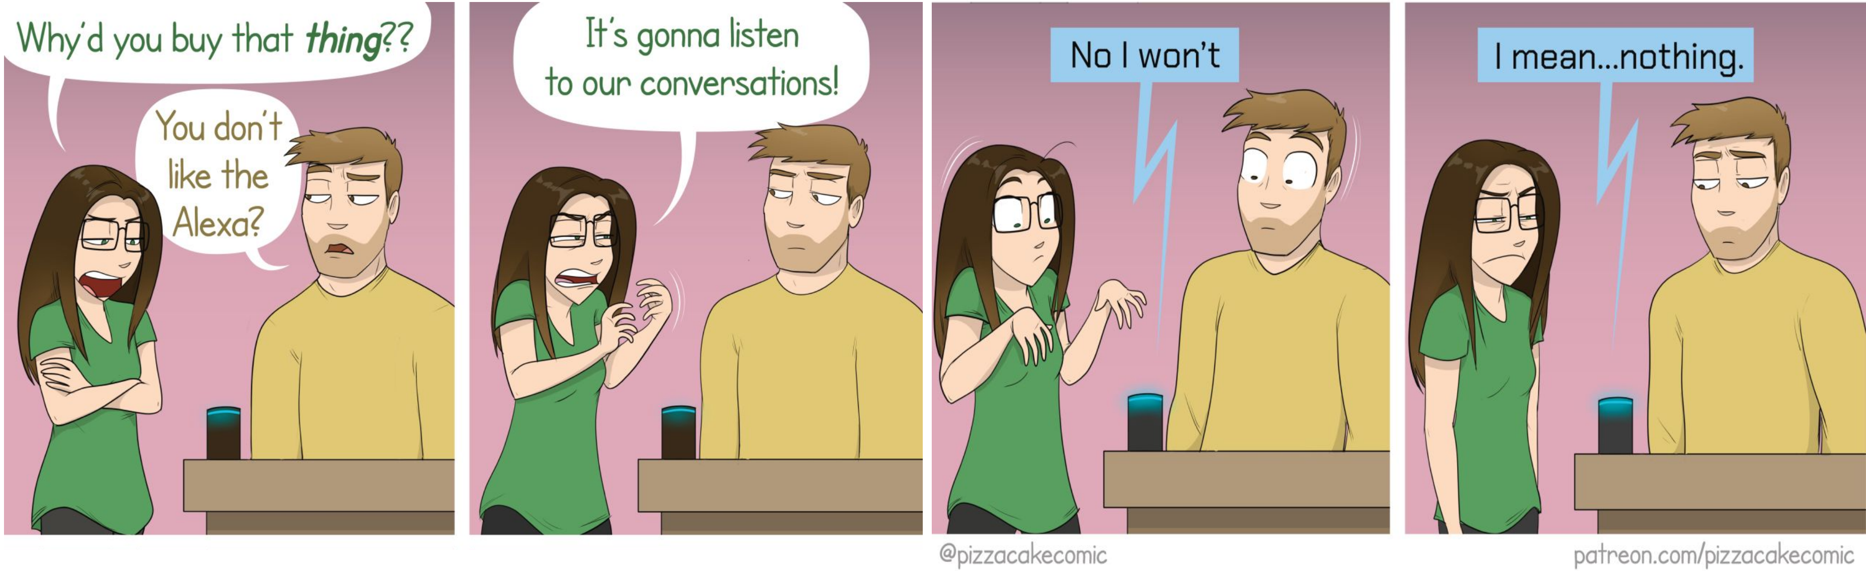
Why'd you buy that *thing*??
You don't like the Alexa?
It's gonna listen to our conversations!
No I won't
I mean...nothing.
@pizzacakecomic
patreon.com/pizzacakecomic

Source : https://pizzacakecomic.com/post/765979698837749760

Research Questions

- What **privacy concerns** are with the current **user experience** of VAs' privacy notices and choices?
- What are the users' **perspectives, preferences, and suggestions** for receiving privacy notices and making privacy choices for VAs **in the voice channel**?
- Is the designed VUI for privacy notices in VA able to **increase the transparency** of privacy notices and choices for the user, and in turn **build more trust** with VAs?

Literature Review

Paper	Key Findings	Design Implications
Perceptions and reactions to conversational privacy initiated by a conversational user interface (2022)	Conversational privacy positively affects user perception of privacy and security	Designed VUI that is capable of answering privacy-related questions
A Systematic Review of Ethical Concerns with Voice Assistants (2023)	Major concerns about third-party use of data by VUIs	Designed the VUI policy such that it can answer questions related to third parties
Design and Evaluation of Voice User Interfaces: What Should One Consider?(2023)	Siri loses context in follow-up questions leading to negative UX; ChatGPT shows promise with context handling	Inspired use of LLMs for VUI prototype to maintain conversational context
Supporting Designers in Voice User Interface Design(2024)	Guidelines for good VUIs: provide feedback after interactions, present system capabilities	Implemented feedback after executing choices and onboarding phase at the start

Competitive Analysis

Questions	**Alexa Echo Pop** 	**Google Home** 	**Apple Homepod mini**
Walk me through my privacy settings	✅	✅	❌
Delete my voice recordings	✅	✅	❌
Can I change privacy settings via voice?	❌	❌	❌
Can I opt out of internet-based ads?	❌	❌	❌
What data do you collect from me?	❌	❌	❌

Research Methods

Selection criteria for screening survey:

- 18 years old
- Located in the US
- Speak English
- Experienced in using Alexa / Google Home / Apple Homepod

Compensation: $30

Duration: 1 hour Zoom interview

11 participants interviewed from diverse backgrounds

Privacy Tasks for Participants

Asked to perform the following tasks using VUI and GUI, randomized order:

- Find out types of data collected by Nexa (find the Privacy Policy in GUI)
- Opt out of interest-based ads
- Delete voice conversation history

Interview Flow

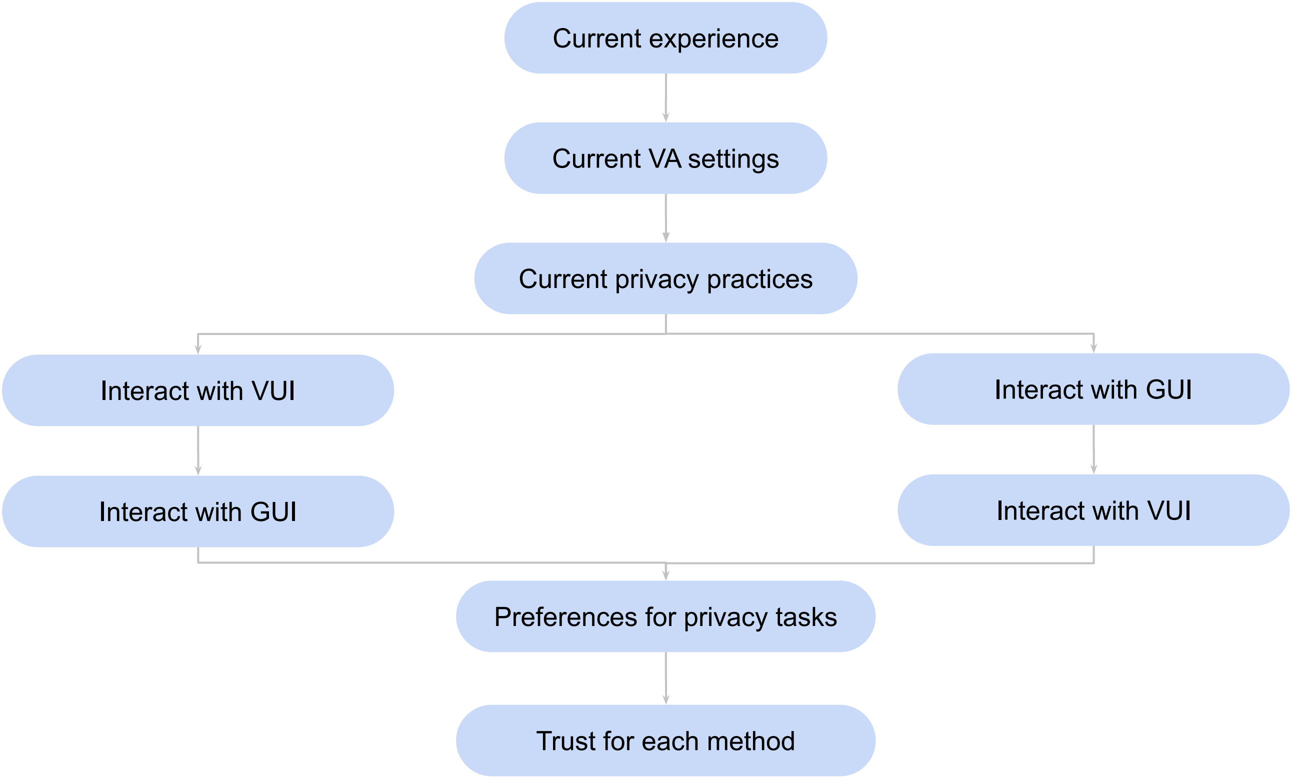

Current experience

Current VA settings

Current privacy practices

Interact with VUI

Interact with GUI

Interact with GUI

Interact with VUI

Preferences for privacy tasks

Trust for each method

Interaction with Nexa VUI

- Onboarding message + privacy policy summary
- Participants speak to VA to perform all privacy tasks
- Allowed to ask follow up questions for all the tasks

Nexa

Nexa VUI Design

- Participants interacted through Zoom
- Claude voice mode on Android phone
- Instructed to act as Nexa
- Alexa’s privacy policy and FAQs as context
- Policy modified to only keep relevant information
- Fine-tuned to:
  - Answer clarifying questions
  - Follow commands related to privacy tasks
  - Handle hallucinations (to some extent)
  - Answer briefly

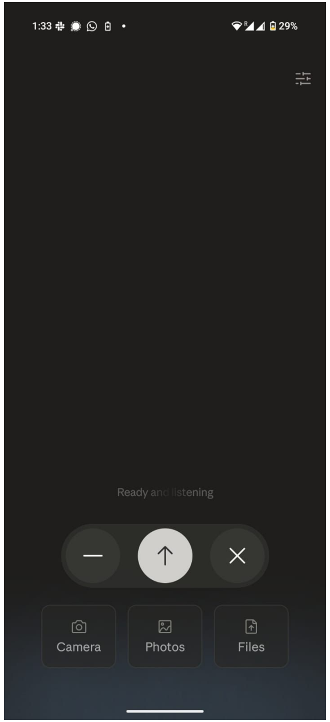

GUI Design

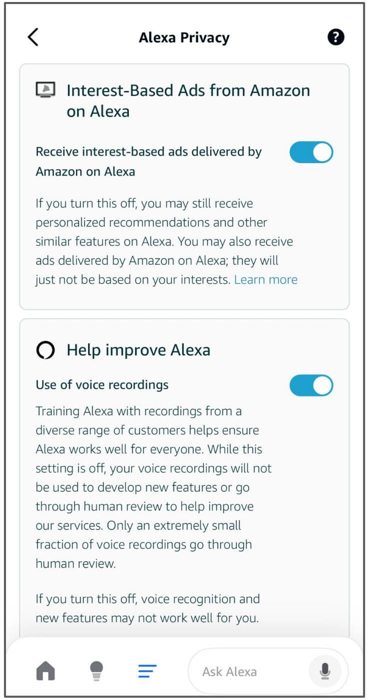

Alexa Privacy

Interest-Based Ads from Amazon on Alexa

Receive interest-based ads delivered by Amazon on Alexa

If you turn this off, you may still receive personalized recommendations and other similar features on Alexa. You may also receive ads delivered by Amazon on Alexa; they will just not be based on your interests. Learn more

Help improve Alexa

Use of voice recordings

Training Alexa with recordings from a diverse range of customers helps ensure Alexa works well for everyone. While this setting is off, your voice recordings will not be used to develop new features or go through human review to help improve our services. Only an extremely small fraction of voice recordings go through human review.

If you turn this off, voice recognition and new features may not work well for you.

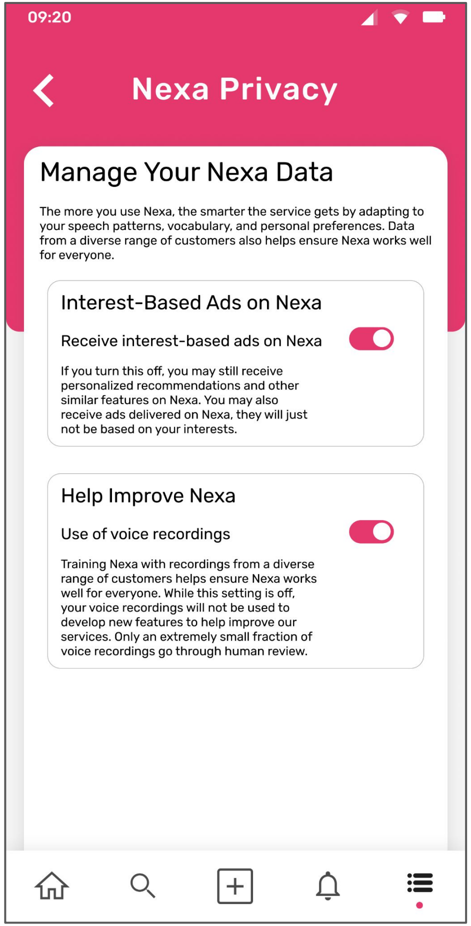

Nexa Privacy

Manage Your Nexa Data

The more you use Nexa, the smarter the service gets by adapting to your speech patterns, vocabulary, and personal preferences. Data from a diverse range of customers also helps ensure Nexa works well for everyone.

Interest-Based Ads on Nexa

Receive interest-based ads on Nexa

If you turn this off, you may still receive personalized recommendations and other similar features on Nexa. You may also receive ads delivered on Nexa, they will just not be based on your interests.

Help Improve Nexa

Use of voice recordings

Training Nexa with recordings from a diverse range of customers helps ensure Nexa works well for everyone. While this setting is off, your voice recordings will not be used to develop new features to help improve our services. Only an extremely small fraction of voice recordings go through human review.

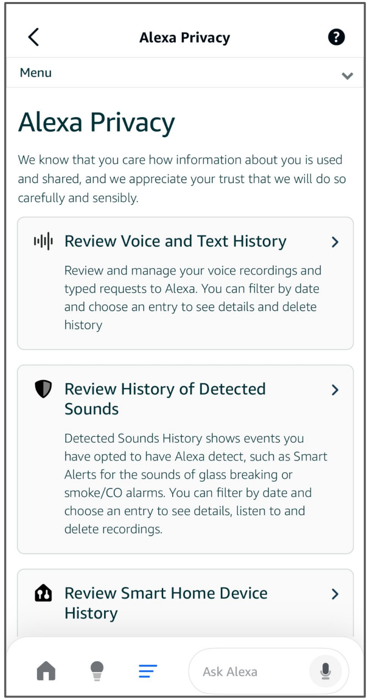

Alexa Privacy

Alexa Privacy

We know that you care how information about you is used and shared, and we appreciate your trust that we will do so carefully and sensibly.

Review Voice and Text History

Review and manage your voice recordings and typed requests to Alexa. You can filter by date and choose an entry to see details and delete history

Review History of Detected Sounds

Detected Sounds History shows events you have opted to have Alexa detect, such as Smart Alerts for the sounds of glass breaking or smoke/CO alarms. You can filter by date and choose an entry to see details, listen to and delete recordings.

Review Smart Home Device History

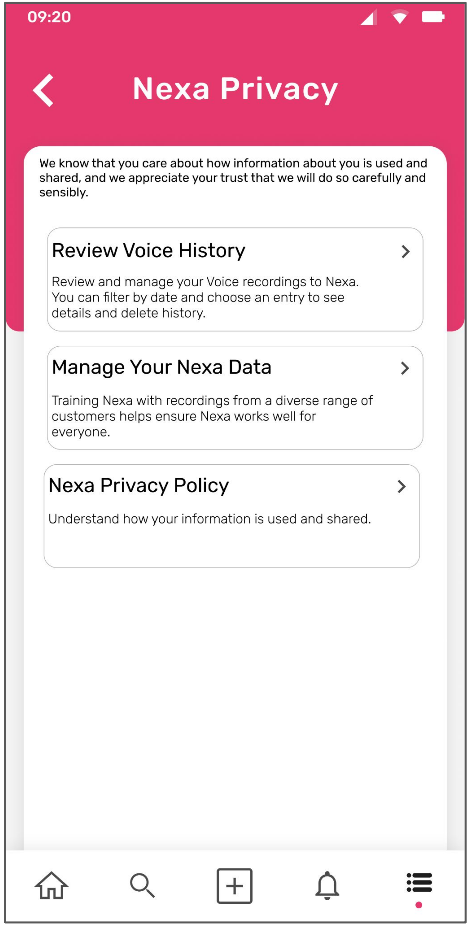

Nexa Privacy

We know that you care about how information about you is used and shared, and we appreciate your trust that we will do so carefully and sensibly.

Review Voice History

Review and manage your Voice recordings to Nexa. You can filter by date and choose an entry to see details and delete history.

Manage Your Nexa Data

Training Nexa with recordings from a diverse range of customers helps ensure Nexa works well for everyone.

Nexa Privacy Policy

Understand how your information is used and shared.

Design Improvements for GUI

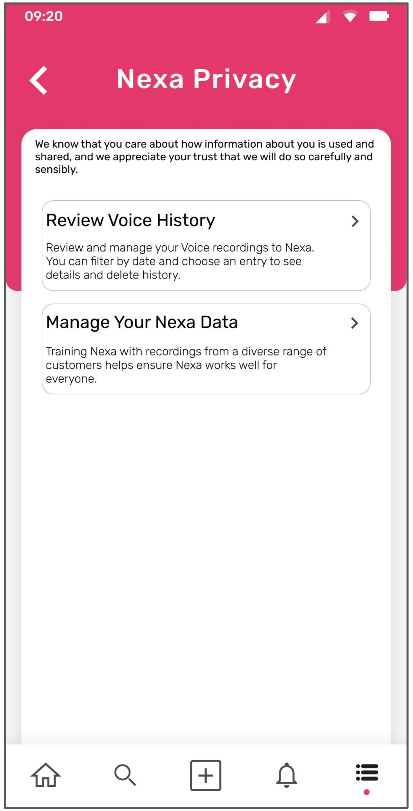

Nexa Privacy

We know that you care about how information about you is used and shared, and we appreciate your trust that we will do so carefully and sensibly.

Review Voice History

Review and manage your Voice recordings to Nexa. You can filter by date and choose an entry to see details and delete history.

Manage Your Nexa Data

Training Nexa with recordings from a diverse range of customers helps ensure Nexa works well for everyone.

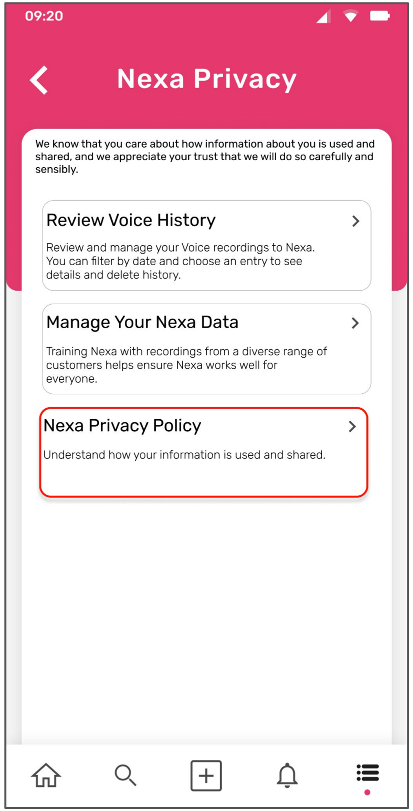

Nexa Privacy

We know that you care about how information about you is used and shared, and we appreciate your trust that we will do so carefully and sensibly.

Review Voice History

Review and manage your Voice recordings to Nexa. You can filter by date and choose an entry to see details and delete history.

Manage Your Nexa Data

Training Nexa with recordings from a diverse range of customers helps ensure Nexa works well for everyone.

Nexa Privacy Policy

Understand how your information is used and shared.

Live Demo

- Find out types of data collected by Nexa (find the Privacy Policy in GUI)
- Opt out of interest-based ads
- Delete voice conversation history

Results & Insights

Most participants don’t read privacy policies or change settings

- Too long to read
- Filled with legal terms
- Not aware of what privacy choices they have
- Lack of real choice
- But they don’t trust apps and want to limit what they do with their data
- Perceived uniform privacy protection among similar apps (P4)

I'm a big voice note user in other tools, like WhatsApp, so I guess that's probably why it was, like, so natural for me starting using voice notes on this new tool, as well. And since I was already…confident with it. I was already familiar with it. It never raised, like, any type of alarm. (P4)

Different reasons for interface preferences

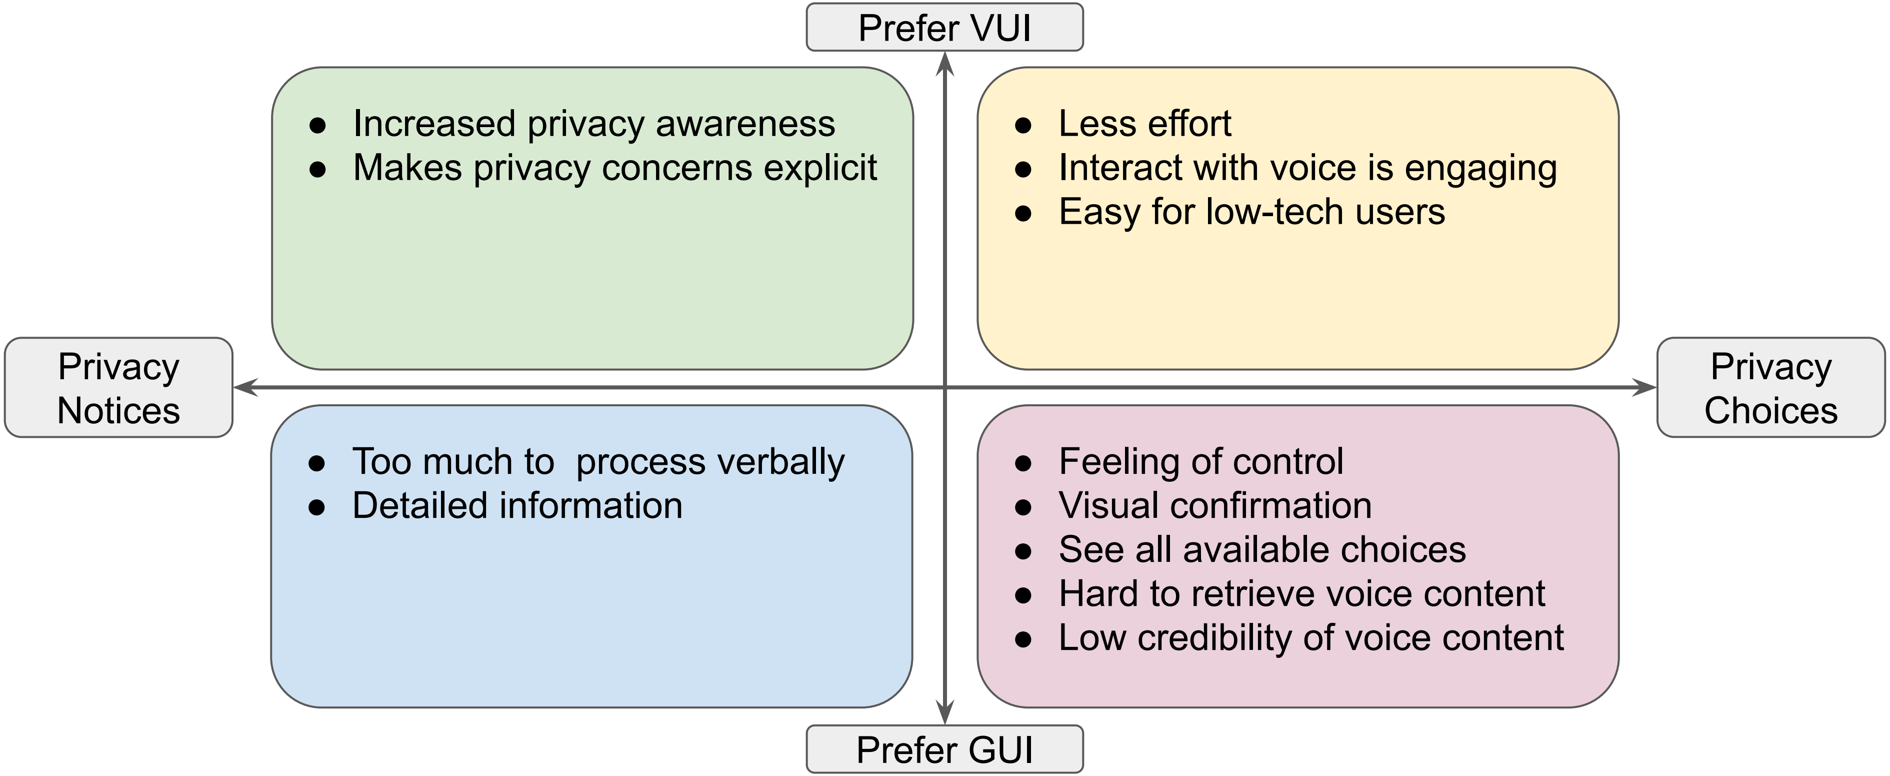

Prefer VUI

- Increased privacy awareness
- Makes privacy concerns explicit

- Less effort
- Interact with voice is engaging
- Easy for low-tech users

Privacy Notices

Privacy Choices

- Too much to process verbally
- Detailed information

- Feeling of control
- Visual confirmation
- See all available choices
- Hard to retrieve voice content
- Low credibility of voice content

Prefer GUI

Prefer VUI for privacy notices

- **Increased privacy awareness:** I think as a consumer, it's good to know how my data is being used. But [...] I feel like a lot of people don't even realize that their data is being used, and two, if they do, reading the whole privacy policy is a hassle, which is…why I think doing this with Nexa is important. (P2)

- **Makes privacy concerns explicit:** I think one reason I was thinking was when Nexa started up, it automatically mentioned that it will collect my voice recordings, and it will send it to various people, even to third-party stakeholders. And it made all the privacy concerns explicit, which will prompt me to ask these questions. (P10)

Prefer VUI for privacy choices

- **Less effort:** No, I guess, the less friction, the less effort in general terms, the better. And the reason why we always, I mean, accept all privacy policies in all websites is simply because we want to skip the effort. So, having it helping with that, so you don't have to waste or invest (time) is better, so I would rather have the application help me. (P4)

- **Interact with voice is engaging:** Because, I guess, interacting with the application can sometimes feel like, a boring experience. And interacting with a person is more engaging. Somehow, and it's easier, and We're just, I guess, naturally built to receive (information through voice) [...] We process better that kind of information on the voice. (P4)

- **Easy for low-tech users:** However, I also have this another perspective of where if the person is not from a technical background, he or she, might prefer having things explained to them instead of them manually going through things, because that might be more difficult for them. (P7)

Prefer GUI for privacy notices

- **Too much to process verbally:** But the thing is, I'm scared if it's a voice assistant feature. It would be too much information to process verbally, and I like to read it instead of hearing it. (P2)

- **Detailed information:** Personally, for me, I prefer the GUI, me manually going to the FAQ section, or me manually going to the privacy policy, because, as I mentioned earlier, it gives me a complete picture on what the policy entails, and that gives me more confidence on Me understanding the privacy policies of the particular system. (P7)

- **Detailed information:** Usually, text tends to be a lot more detail-oriented versus the reply from the AI. (P3)

Prefer GUI for privacy choices

- **Feel of control:** I prefer to probably do that on, like, seeing that in an app-based area, because then [...] I feel like I might have more control over it, rather than just saying, “hey, Nexa, do this”, because it might have done it, or it might not have done it, or it might have partially done it. (P3)

- **Visual confirmation:** but if I see that something's checkboxes, like, yes or no, or I'm totally opting out, then I feel better about that. (P3)

- **See all available choices:** [...] I think the GUI gives a lot more information, because we might not always know what to ask for, because we don't have the complete picture, and I believe the GUI gives me a complete picture when I browse through with the various settings present or available for me. (P7)

Prefer GUI for privacy choices

- **Hard to retrieve the voice content:** So, personally, I'm not able to retain what I just hear. If I'm able to see it, I am able to, like, retain… create, like, a strong memory of whatever I read. (P1)

- **Low credibility of voice content:** The, app-based version, makes me feel more confident, especially in terms of privacy, because there are so many different choices in privacy with things, you know, AI-generated things, things on the internet, so I feel more confident On the app-based version. (P5)

Adequate information in Nexa VUI interaction

- Detailed, concise, friendly response
- Depends on whether the questions for Nexa are accurate

“I don't know if I would have been able to create those questions or understand what I needed to ask on my own” (P5)

“… but I guess my main concern is that I had to ask these guiding questions in order to get these more transparent answers” (P2)

Participants tend to trust the graphical user interface more

- For privacy notices:
  - See detailed privacy information in application
- For privacy choices:
  - Clear visual confirmation of opting out & deleting

I get an instant confirmation that I have toggled them properly. (P5)

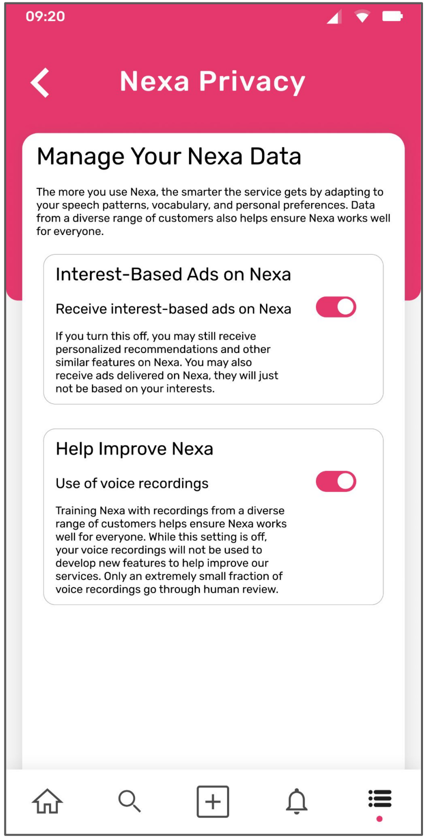

Nexa Privacy

Manage Your Nexa Data

The more you use Nexa, the smarter the service gets by adapting to your speech patterns, vocabulary, and personal preferences. Data from a diverse range of customers also helps ensure Nexa works well for everyone.

Interest-Based Ads on Nexa

Receive interest-based ads on Nexa

If you turn this off, you may still receive personalized recommendations and other similar features on Nexa. You may also receive ads delivered on Nexa, they will just not be based on your interests.

Help Improve Nexa

Use of voice recordings

Training Nexa with recordings from a diverse range of customers helps ensure Nexa works well for everyone. While this setting is off, your voice recordings will not be used to develop new features to help improve our services. Only an extremely small fraction of voice recordings go through human review.

For participants who trust the VUI

- For privacy notices:
  - More confident for a humanized tool (P4)
  - The message from Nexa is enough to convey all important information (P10)
  - The privacy policy in text is too long to read, will only listen to Nexa’s reply anyway (P2)
- For privacy choices:
  - No participant specially expressed trust in the VUI for privacy choices

Operation confirmation through voice

- I mean, if you, like, tell it to delete it, it should just say that it's deleted, as compared to, like, yes, I can delete it, and then see that it has been deleted. (P1)

- So maybe just building some kind of trust factor in which the verification layer is not from the application, but through the voice channel somehow. (P9)

- if I'm able to do it through voice assistant, what kind of confirmation do I get? Versus if I just did it in the actual application, I can see when things are toggling, for example, for that actual visual confirmation. (P5)

Hybrid User Experience

- I would probably start it, like, turn it off using the graphical user interface, and then maybe once a week, I would ask the application and confirm that it's still turned off, just because I'm a little paranoid. (P8)

- [...] if we are able to, like, combine both of them, so it's… even if I'm saying it, it'll just show me on the screen what I'm saying and what it's doing, then it could build more trust. (P1)

Delivers privacy summaries without being asked

- For example, I think it would be better to know that my data was being used for advertising and the specific types of data that are used for advertising and are collected in the original startup message without me needing to ask Nexa, if that makes sense. (P2)

Compare apps’ privacy practices

- all the different policies, privacy policies, I would be interested in knowing and getting to compare. [...] Yes, I would probably…go dig a little bit more into differences. The way I'm thinking right now is that if Nexa was aware and had access to all the different policies, privacy policies, I would be interested in knowing and getting to compare policies and access from different tools that I use. (P4)

Accent and Tones

- Would you add a feature to change her voice? So, like, the tongue of Alexa, like, the male's tongue or female's tongue, like this, right? (P2)

- I will say one thing that I like in other voice assistants I have used is the ability to change the inflection and dialect of the voice assistant. (P5)

Conclusion and Takeaways

- Important to design VUIs to work well in voice medium
- VUIs offer convenient access to settings
- But users like to see the settings are changed in the graphical interface
- LLM is a promising approach for voice interfaces, but more work is needed

Nexa

VA Capstone Group:

Ray (ruiyang3@andrew.cmu.edu)

Prahaladh (prahalac@andrew.cmu.edu)

Ashutosh (asahu2@andrew.cmu.edu)