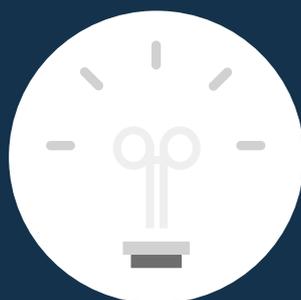


Transparency & Control in Conversational UX

Design considerations for
AI-powered voice interactions



Introduction

We are witnessing a fundamental shift in how humans interact with technology, as voice-based interactions become not only common, but often the primary user experience. From voice notes to smart devices and AI assistants.

Over the last twenty years Graphical User Interfaces (GUIs) have been dominant, but we are now seeing the rapid emergence of Conversational User Interfaces (CUIs) enabled by Large Language Models (LLMs). These new AI-powered conversational experiences are transforming the potential of voice and natural language as an interaction method, creating an open-ended, adaptive, and conversational user experience.

This shift moves beyond the "visual templates" model of the past to a new era of responsive, context-driven interactions with apps and services. We have new opportunities to improve the user experience, build on all that we've learned over the past few decades of designing Transparency & Control moments in visual GUIs, and explore new opportunities and challenges in adapting those experiences through the addition of a conversational, voice-first approach.

Authors

Lead



Peter Tanham
POLICY MANAGER
(LEAD AUTHOR)



Adam Bargroff
POLICY MANAGER



Dan Hayden
POLICY DIRECTOR



Andrew Evans
CONTENT DESIGNER



Shari Husain
DIRECTOR,
CONTENT DESIGN



Vivek Mayasandra
PRODUCT DESIGNER



Shane Mannis
USER EXPERIENCE
RESEARCHER



Sara Bastomski
USER EXPERIENCE
RESEARCHER



Tamy Guberek
USER EXPERIENCE
RESEARCHER



Elaine Montgomery
DIRECTOR,
PRODUCT DESIGN

How to cite this report?

Peter Tanham, Dan Hayden, Adam Bargroff
"Transparency & Control in Conversational UX",
(March 2026) at <https://ttclabs.net/reports/>

Thank you to our design partners at **CraigWalker Design and Research**, who helped us transform our research into a well-designed report.

What you'll find in this report

1	Overview	5
2	Opportunities for Transparency & Control within conversational UX	9
	Emerging experiences enabled by AI	10
	Empowering natural dialogue	12
	Elevating personalized control	13
	Unlocking deeper engagement	13
3	A new toolkit	15
	Emerging considerations	16
	Questions to consider	17
	Encountering a Transparency & Control moment	18
	What is an appropriate user experience?	18
	What considerations should we take into account for the style and tone of content?	19
	Sequencing information	21
	How do we prioritise information in a linear voice-based experience, while maintaining a good user experience?	21
	Multimodality	24
	Do we envisage Transparency & Control being possible via voice only, or will we also need a screen?	24
4	Summary	27

1

Overview

About this report

This report presents the feedback and discussions from a series of global workshops exploring these new opportunities and challenges. The workshops engaged academics, civil society, industry, legal professionals and other policy stakeholders in Berlin, Madrid, Seoul, Brasília, San Francisco, and Delhi.

Participants engaged in Design Jam workshops to evaluate scenarios across a range of Transparency & Control moments, with a focus on voice as the input, to what a great user experience might look like in these moments.

The insights from these Design Jams are shared here as early insights into this area. As the program evolves, we intend to explore and publish additional insights on this topic.

This report is about	This report is not about
✓ AI-powered interaction with voice as the primary modality	✗ A comprehensive coverage of all Conversational UX, which might also include considerations around screen and text-based UIs.
✓ Voice as an input method where the user's instructions, requests and responses are the UI	✗ Processing voice as data, for example, to identify a user or analyze voice patterns
✓ A selection of interesting and relevant Transparency & Control moments	✗ The comprehensive toolkit of Transparency & Control required to support AI-powered Voice or the broader Conversational UX
✓ An initial exploratory engagement with academics, policy makers, product designers and legal professionals	✗ A definitive set of rules or recommendations

How to use this report

The following exploratory examples for AI-powered voice Transparency & Control moments have been developed to support product makers and policymakers in their emerging thinking around designing Transparency & Control in Conversational UX.

The aim of this report is to capture and share initial insights and recommendations, alongside key examples and questions to stimulate further discussion and exploration; as such, the report is not intended to be definitive or to provide a series of specific guidelines or standards.

The following sections of this report will explore:

- **Key insights and an emerging toolkit** around the considerations highlighted from this early engagement

- **Key recommendations** for Transparency & Control moments of AI-powered voice

Glossary of terms

Transparency

In this report, Transparency refers specifically to Data Transparency - information presented or available to users about how their data is collected, stored and used.

Control

In this report, Control refers specifically to Data Controls - settings that users can change to affect how their data is processed. This can include “on demand” controls, like Settings, or “up front” controls, like a consent.

T&C Moment

A point in the user journey where Transparency is shared or a Control is surfaced.

Conversational UX

User experience designed around natural language interactions, often using voice or chat interfaces.

Voice User Interface (VUI)

An interface that allows users to interact with technology using spoken commands and responses.

GUI (Graphical User Interface)

A visual interface with elements like menus, icons, and buttons for user interaction.

Large Language Model (LLM)

An AI system trained on vast text data to understand and generate human-like language.

Rule-based VUI

Early voice systems that rely on fixed commands and scripts, lacking flexibility or context awareness.

AI-powered VUI

Voice interfaces enhanced by AI, capable of understanding context and open-ended requests.

Progressive Disclosure

A design approach where only essential information is presented first, with further details available on demand.

Adaptive Disclosure

Adjusting the amount and type of information shared based on user context or preferences.

Linear Experience

An interaction where information is presented in a fixed, sequential order, typical of voice interfaces.

Multimodality

The use of multiple input/output channels (e.g., voice and screen) in a user experience.

Earcon

A brief, distinctive sound used to signal a specific event or state in a voice interface.

2

Opportunities for Transparency & Control within conversational UX

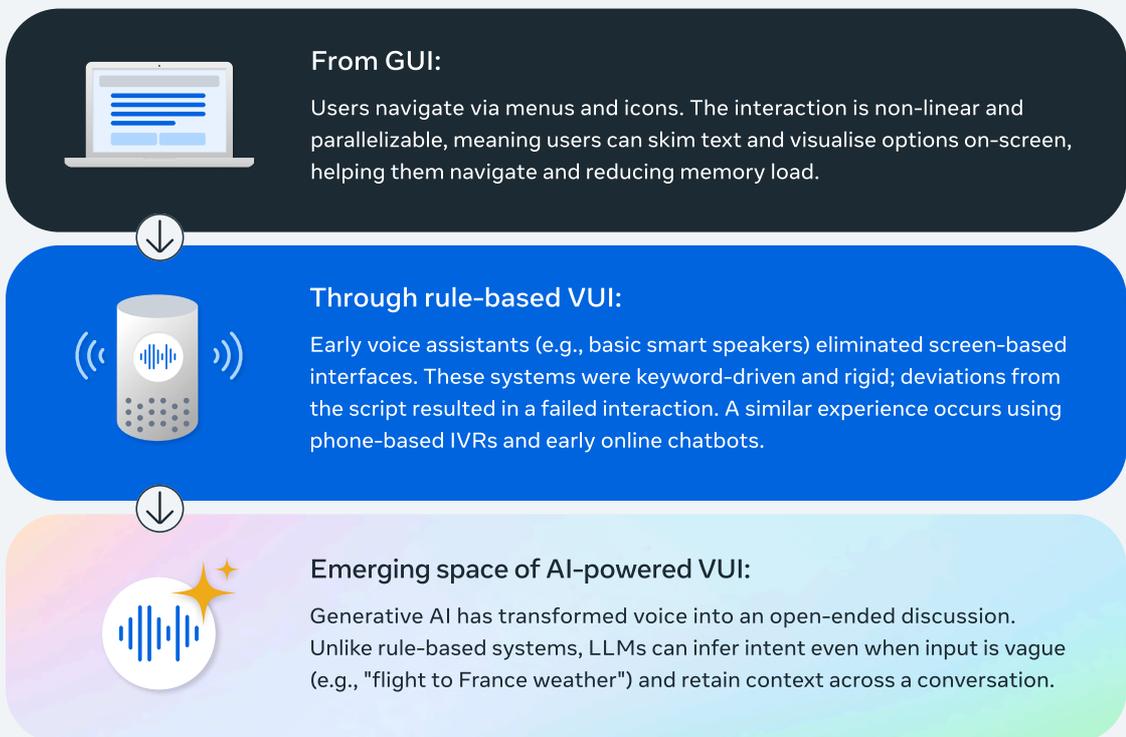


Emerging experiences enabled by AI

The user experience is evolving as new services, powered by AI, shift users from explicit visual controls on a graphical interface to conversational and voice-based interaction models. This is creating a paradigm shift in how users interact with technology and how we present Transparency & Control moments to them.

This report draws on our early explorations of a more specific component of the conversational UX space: the use of voice as both input and output, a context in which the user is unlikely to rely on a screen.

Over the last twenty years, we have seen a slow progression from Graphical User Interfaces (GUIs) to Rules-Based Voice User Interfaces (VUIs) and very recently to VUIs powered by Large Language Models (LLMs).



This paradigm shift means we are moving from a screen-based interface, in which users were steered by the interface, and increasingly to voice, in which users expect a more conversational, responsive experience. This means we are seeing the user experience go:

From	To
Pre-defined choices	Open-ended requests
Visual templates	Conversation and discussion
Predictable flows	Adaptive, context-driven reasoning

It expands a user's perception of options while reducing the interface's ability to shape that perception.

This shift means we cannot simply port across the GUI design patterns we currently use for Transparency & Control, e.g., cookie banners, options menus, buttons and toggles, and privacy settings pages. In this new channel, Transparency & Control must be spoken, contextual, and dynamically generated within a user flow.

The design space for this new AI-powered voice experience should support conversational, adaptive disclosures for how people interact with an AI-powered VUI, rather than disrupting conversational flow.

An example might be an AI stopping to read a disclaimer in the middle of a sentence, creating a jarring and awkward conversation.

An example might be during onboarding: reading a full policy aloud is a poor user experience and is likely to cause users to tune out or abandon the task because it feels too long.

When viewed across the end-to-end product lifecycle, we can begin to understand not only the opportunities this presents but also the challenges we face as we evolve the Transparency & Control user experience.

	Opportunities	Challenges
Upfront	I can interrupt, ask questions, dig deeper and move at my own pace	A linear onboarding experience could be too long to sit through
In-context	Not just the right time, but also the right tone, language and information	In-context messages and permissions can interrupt my experience
On demand	I can ask an AI exactly what I want, or explain my goal, without navigating settings menus or help centres	Understanding what I could ask is hard, as there is nothing visual to guide me

[Using this Transparency & Control model, we can see how voice changes the experience for both the user and the provider]



Empowering natural dialogue

There is an opportunity to improve accessibility by embracing a natural, conversational style of communication.

By leveraging AI, we can create experiences that feel more accessible and colloquial. This ability to engage in a "having a chat" manner helps overcome the feeling of "excessive legalese" that users often report, or the need to learn and endure a rigid system ("press 9, then 0, then # to speak to an operator" and generally makes Transparency & Control more accessible to everyone, regardless of their technical literacy.

EXAMPLE CHAT

Can you make a restaurant reservation for me? Tomorrow at 12pm at The Big Burger Restaurant on Main Street.

To make a reservation for you, I'll need to share your name, email, and phone number with the restaurant —is that okay?

Yeah, sure.

During our workshops, we discussed that allowing users to provide consent through everyday language, such as "fine" or "okay", creates a much more seamless, human-centric flow.

Policy stakeholders did share concerns that current technology may have limitations and may not parse affirmations reliably. "What if the user says 'fine' sarcastically? What if they say 'huh?' but the system hears 'yeah'".

In these cases, the direction from our policy stakeholders was to accept a range of affirmative and negative terms, but to implement fallbacks and safeguards where there is any ambiguity (such as asking the user to clarify). Long-term, as it was felt that the AI could parse affirmation with a high level of reliability, we could move to a more open-ended dialogue for Transparency & Control responses.

Colloquial Transparency also holds promise to help users get past the "wall of legalese text" problem we often face in visual interfaces, allowing them to ask questions and gain understanding in a personalized, accessible way. Most policy stakeholders also expressed a positive attitude to this approach, with the caveat that the information should remain accurate. Some stakeholders maintained a preference for the more formal language to remain consistent across all users.

While companies refine AI reliability to respond to a wide range of prompts, they should seek to deliver an experience that is both a legally robust, formal engagement and a much more naturally intuitive conversational UX.



Elevating personalized control

Transitioning to a voice-first, AI-driven UX opens the door to highly personalized Transparency & Control.

We have the chance to move away from rigid "one-size-fits-all" systems of notices, controls, and consents toward ones that understand and respond to the specific context of a user's life.

While many settings and control interfaces aren't intentionally complex, they often become difficult for users to navigate because they must accommodate all the options a user might need. This can lead to a user's preferred controls being buried within settings menus, making it hard for users to understand and choose the settings they require.

Conversational AI bypasses this structure by simply allowing the user to describe their desired outcome. This shifts the burden of navigation from the human to the system, transforming a multi-step search into a single, voice-driven request.

Many Policy stakeholders were excited by this prospect of more personalized controls, with some suggesting an even more granular approach to consent:

“Yes, share my location for the next 3 days while I'm in New York.”

This level of granular, conditional consent has the potential to increase user agency and engagement with personal data use and privacy. As an opportunity to evolve from static interfaces to dynamic, language-based preference management, the AI acts as an intelligent translator, recognizing the request and proactively following up to ensure the user's intent is fully captured and executed.

While it was recognized among policy stakeholders that we were likely some way off from being able to respond at such a granular level today to a narrowing of the scope of consent, we can mitigate early issues by proactively following up with additional questions to clarify the user's intent.



Unlocking deeper engagement

Shifting to a conversational UX isn't just about convenience; it can also help users build deeper trust in the products and services they use.

By replacing static FAQ pages with a free-form, user-led Q&A, we enable users to explore Transparency & Control moments in ways that are more relevant to them, in a language and style that is more engaging.

This approach can transform transparency and control from a step in the user experience into a helpful conversation that helps users pinpoint exactly how their data is used in specific, real-world scenarios and deconstruct complex policies using everyday language.

There is also the potential for AI to leverage historical context from the user's previous interactions, further supporting their education in how to personalize their privacy goals and knowledge.

EXAMPLE CHAT

“What did we do last time I went on a trip overseas?”



Last time you allowed this for a similar booking app, would you like to do that again?

This experience moves even closer to a tailored position for an individual, from a simple set of controls and consents to a highly personalized experience that fosters a sense of being proactively understood and protected, rather than merely presented with a set of toggles.

The next section of this report focuses on key Transparency & Control moments identified as early issues in an AI-powered voice user experience and examines the new toolkit we require.

3

A new toolkit

Initial design considerations
as we move to a voice-based
conversational UX

Emerging considerations

As we transition from graphical interfaces to AI-powered voice, we encounter a unique set of tensions. We rely less on visual cues like bold text or font size to signal importance. Instead, we need to find new ways to communicate transparency and control without breaking the natural flow of conversation.

The GUI era of Transparency & Control relies on users' ability to visually scan, access parallel information, and receive information as persistent text. In voice, information is transient (once said, it is gone) and potentially linear (you cannot "skim" a spoken privacy policy). This means that our current toolkit of GUI Transparency & Control design patterns is largely unavailable to us if we focus primarily on voice and do not rely on the screen as a key source of information.

We cannot simply recreate our toolkit of GUI controls in AI-powered voice experiences; therefore, we will need to review and potentially build new experiences that preserve user agency when the information shared is temporary, sequential, and dynamically generated.

Toolkit	GUI today	AI-powered voice
Layering	<p>Layered information (The "Learn More" Pattern)</p> <p>A GUI allows a user to drill down into details and immediately snap back to the main context.</p>	<p>In voice UX, navigating away from the primary consent flow to hear a long, detailed policy, then trying to "return" to the specific decision-making point, can be disorienting and disrupt the conversational flow.</p>
Toggles	<p>Selective consent with granular toggles</p> <p>In a visual interface, a user can scan a list of five different permissions and quickly toggle the two they dislike.</p>	<p>In a voice environment, the system would need to read every option aloud sequentially and pause for a "yes/no" response to each, causing significant friction, cognitive load, and a poor user experience.</p>
Tables	<p>Tiered consent tables</p> <p>Visually comparing features across tiers allows users to "weigh" options simultaneously.</p>	<p>Replicating this via voice requires the user to hold complex comparative information in working memory, which could be difficult and hard to remember without a visual reference.</p>
Proximity	<p>Visual grouping of data types</p> <p>On a screen, visual proximity communicates that items belong together.</p>	<p>In voice, the system must list them linearly, e.g. "We need your location, AND your age, AND your family status...". The serialisation of the "bundle" could feel like a long list of separate demands rather than a single, simple, cohesive request.</p>
Icons & colours	<p>Non-verbal accessibility cues</p> <p>Using icons to support text and contrasting colours on screen to differentiate between affirmative and negative options allows users to scan and understand the sentiment of a button or the nature of a data request without reading every word.</p>	<p>Voice interactions cannot use a "green button" to signal a positive action or an icon to quickly represent "location data".</p> <p>The "scanability" provided by these visual tools is lost, requiring the system to verbally explain every detail that the eye would otherwise absorb instantly.</p>

[A comparison of well-used Transparency & Control design patterns across the two channels]

Questions to consider

This section outlines the primary design considerations identified during our global workshop series with policy stakeholders. These are not intended to be a comprehensive or final set of standards. Rather, they serve as an early exploration of the "known unknowns" - the open questions that teams should ask themselves as they embark on this journey, and the areas where further discussions and exploration could help to define new best practices.

Through our workshops, three core themes emerged as the early foundations for a successful voice-based Transparency & Control moment:



Encountering a Transparency & Control moment

How do we signal to a user that they are entering a T&C moment without visual cues?



Sequencing information

How do we prioritise information in a linear, ephemeral medium?



Multimodality

Can voice be a standalone channel, or must it be supported by a screen?



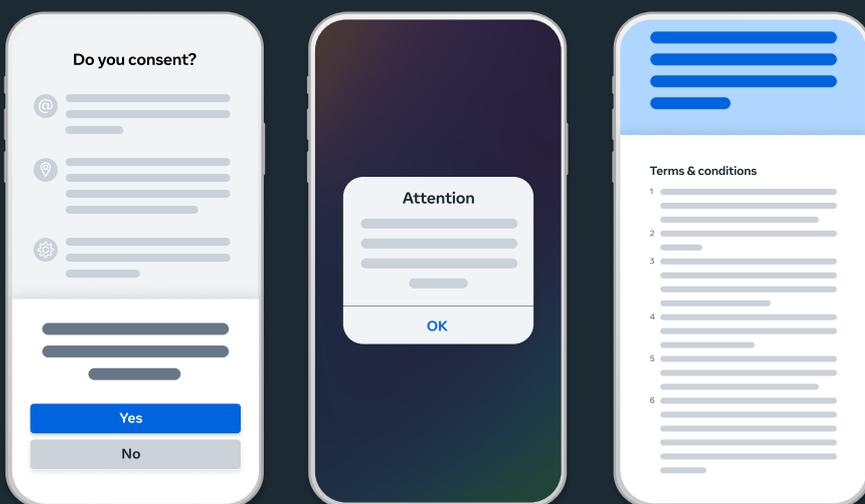
Encountering a Transparency & Control moment

What is an appropriate user experience?

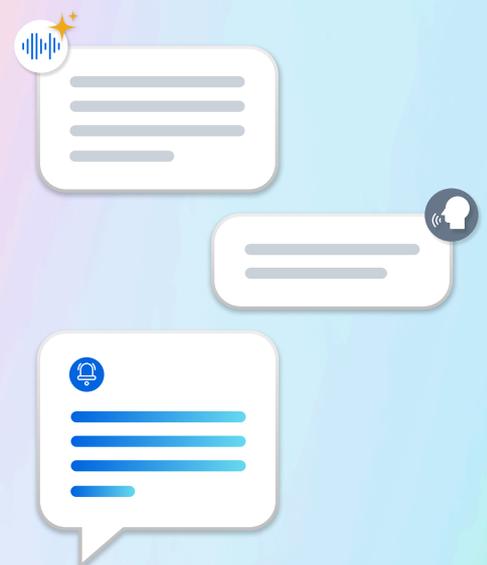
In a GUI, Transparency & Control moments often appear different, highlighting important information and presenting the user with options around making an informed decision. The content is explicitly designed to help users understand the information and quickly scan details, often with links to additional information.

What voice equivalents exist for these GUI design elements?

The bold text of a headline or the dark grey font used for secondary info, a section in italics or on a grey background? To distinguish these moments from the broader product experience, how should we provide similar cues and highlights in voice? And in the era of generative experiences, should they remain consistent across users or adapt to their personal preferences?



From a series of familiar visual templates...



... to a dynamic conversational flow



Signifying a transparency & control moment

An initial consideration of our exploratory workshops was how to signal to the user that they were in a Transparency & Control moment within a broader product experience. What might signal the difference? Would the inclusion of a range of toned or styled differences support the user in knowing: the inclusion of a bell sound effect, or a change in the tone of language? Could we even consider AI pretending to be a character?



Focusing on important content

Across the workshops, it was generally not considered necessary to make a transition from a general product experience to Transparency & Control moments audibly distinct, allowing them to blend with the wider product experience. It was also recognized that the ability to mark critical moments should be used in moderation and non-intrusively, especially in interactions where safety is a concern, such as driving.

Within the T&C moment itself, alerting users to important content and guiding them to focus on it was considered important in voice-based conversations, where earcons (sound cues) could play a supportive role.

“It feels like a cue that helps me refocus right before consenting.”

- KOREA

☆ Recommendation

Earcons (sound cues) are important for alerting users and supporting their focus on voice interactions.

What considerations should we take into account for the style and tone of content?

Another aspect we considered was the tone of voice: whether there is user discretion in its deployment and whether this affects users' perception of the information's formal nature. It sparked discussions about whether the tone of voice should use local slang to improve accessibility or whether it is necessary to maintain a more formal legal language.



Optionality for style

A general consensus from the workshops was that while warm, accessible natural language is welcome to users, concern was raised that if not used appropriately, it could dilute the formality of a Transparency & Consent moment.

Providing users with control over some aspects - such as a more familiar and personable style, or faster talking speed - was likely to enhance their access to critical information. However, it was noted that overly informal language risked downplaying the seriousness of legal decisions, and it was felt that content is the priority and that stylistic aspects, such as tone, matter less.

“They were way too casual with me, as if I were below them. It made them seem unprofessional.”

- KOREA

There may be a middle ground on personalisation that warrants exploration, allowing Transparency & Controls to be delivered in ways accessible to each individual user while maintaining a level of consistency that supports legal confidence.

★ Recommendation

If useful, allow users to control the stylistic aspects of voice experiences, including tone and speaking rate.

EXAMPLE CHAT

To give you a list of movies you might like, I'll need to explain how our feature works and make sure you're good to go. Wanna do that now?

Yeah, sure.

SIGNIFYING A T&C MOMENT

Participants agreed that earcons (sound cues) are important for alerting users and helping them focus on the voice interaction.

I can recommend something for you by analyzing your streaming service profile, which means I'll be able to see what movies you've watched and the ratings you've given them. Should I tell you more or do you want me to start analyzing your streaming service profile?

Tell me more.

OPTIONALITY FOR STYLE

Preferences for the AI's tone of voice varied—some wanted a formal style, while others preferred a warmer, more informal approach.

Ok. If you give me permission to analyze your streaming service profile today, I'll continue to do this in future for every movie you watch. This means I'll be able to give you the most accurate recommendations in future. So now I've read you the information, do you want me to start analyzing your streaming service profile now?

Yes.

Okay, I'll start analyzing the movies you've watched on your streaming service... it looks like there's a romantic comedy and a crime thriller you might enjoy.

Cool, let's go for the crime thriller.



Sequencing information

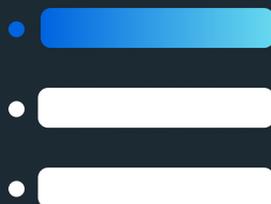
How do we prioritise information in a linear voice-based experience, while maintaining a good user experience?

Voice-based conversation is linear in nature. With a GUI, a user can read and navigate content as they choose. With voice, this is no longer available to the user, so they are reliant on the product maker's decision about how information is sequenced in a Transparency & control moment.

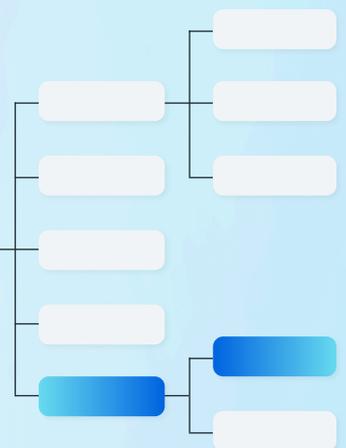
The need to disclose information in bite-sized chunks was seen as important for the user experience. As one policy maker in the workshops noted:

“It’s a better user experience to not make them listen to too much at once.”

- BERLIN



From a screen that helps users scan back and forth ...



...to a non-visual journey where we need to manage the volume of content



Don't try to share all information at once

For disclosing transparency information, progressive disclosure was preferred, as it is easier to consume in a back-and-forth voice interaction than as a single large block of text read aloud. It also gives the user greater agency throughout the experience.

Aligned with a preference for progressive disclosure was the expectation that the system should focus only on information immediately relevant to the current task, deferring nonessential details to a later stage. For services in which users already have a strong expectation regarding data use (e.g., location sharing in ride sharing), the disclosure can be significantly shorter.

It is good practice to let a user know how to change their settings and revoke a consent. In visual consent flows this information can be provided on the same screen as the request - i.e. before the consent is granted. In a voice-first context, policy stakeholders agreed that providing this information after the consent decision is recorded was acceptable.

“Sensitive data, personal information, and third-party sharing are critically important in Korea.”
“When I’m driving, I just don’t have the cognitive capacity. Anything long or repetitive becomes a distraction.”

- KOREA



Skipping non-critical information

"Many times, I just want to skip. It's natural to skip in human conversation."

- BERLIN

"If it keeps talking after I skip, it feels aggressive."

- KOREA

There were also strong expectations that, if critical information (e.g., contractual changes, information with financial implications or of a sensitive nature) is shared upfront and succinctly, users should not be permitted to skip this information, as it could create legal risk.

However, it was recognized that presenting every detail up front in a voice-first experience could be too burdensome for the user. Therefore, it was felt that it would be better to allow users to skip non-critical information, as is likely their expectation.

"[While ideal to offer the information upfront, it offers] a better user experience not to make them listen to too much at once."

- BERLIN



Overall recommendations

- Use of progressive disclosure to avoid overwhelming the user.
- Allow users to skip non-critical information during a T&C moment.
- Provide information on revoking consent after the decision is made.

EXAMPLE CHAT

 To give you a list of movies you might like, I'll need to explain how our feature works and make sure you're good to go. Wanna do that now?

Yeah, sure. 

 I can recommend something for you by analyzing your streaming service profile, which means I'll be able to see what movies you've watched and the ratings you've given them. Should I tell you more or do you want me to start analyzing your streaming service profile?

Tell me more. 

PROGRESSIVE DISCLOSURE

Participants consistently expressed a preference for progressive disclosure (where information is given in smaller, interactive chunks) as it gives users more agency and is easier to process.

 Ok. If you give me permission to analyze your streaming service profile today, I'll continue to do this in future for every movie you watch. This means I'll be able to give you the most accurate recommendations in future—

—Yes, okay, skip, just give me some movie ideas. 

SKIPPING

Participants agreed that skipping can be allowed for low-risk, non-contractual information. This was contingent on the expectation that users are clearly informed about what they're skipping and can easily retrieve the information later.

 Okay, you can ask me about permissions any time in the future. Now I'll start analyzing the movies you've watched on your streaming service... it looks like there's a romantic comedy and a crime thriller you might enjoy.

Cool, let's go for the crime thriller. 

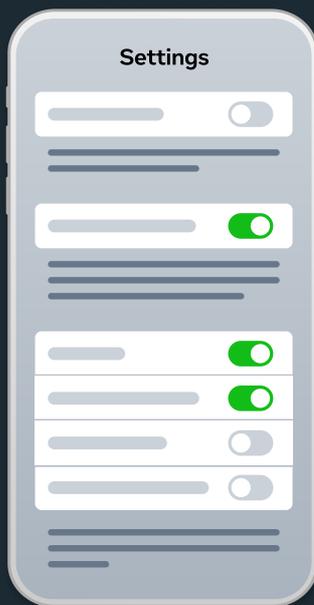


Multimodality

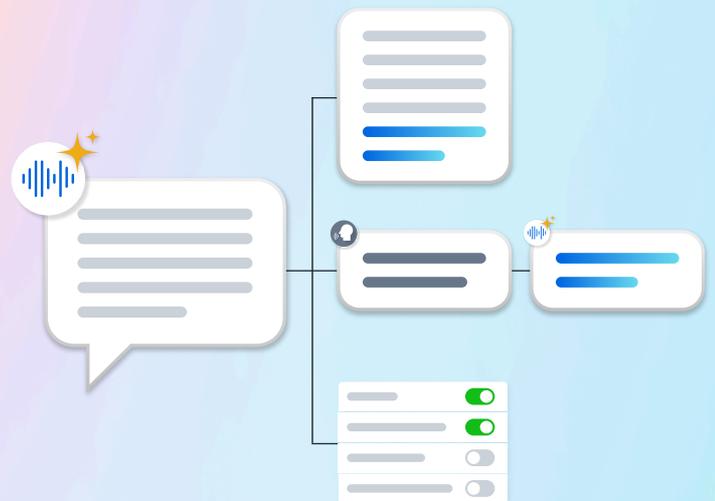
Do we envisage Transparency & Control being possible via voice only, or will we also need a screen?

As our devices evolve to become wearables, like AI-powered glasses, ear buds and smartwatches, and as voice AI becomes more accurate and powerful, we expect that voice will become the primary modality in some or even most circumstances. In these cases, what will the role of the screen be? An optional companion? A necessary hand-off? Or something in between.

Can all of Transparency & Control be conducted through voice, or should some still happen on a screen? In this section we explore the ways in which an additional device can help users manage and navigate these experiences.



Moving from menus, toggles and settings to manage preferences...



... to using voice commands



A visual dashboard to review

A centralized on-screen control centre or dashboard, through which users could review transparency settings and decisions, was perceived as a positive and supportive experience. Because voice interactions are temporary, this was considered a necessary feature and likely one that users would expect.

The consensus was that this central "control centre" would enable users to review and adjust privacy settings and the decisions they had made via voice.

"Psychologically, it's helpful to have a designated space [to confirm settings/decisions]".

- BERLIN



A transcript for clarity

Another area discussed around managing Transparency & Control settings in voice was the use of screen-based summaries. This was regarded as beneficial for both the user and the business. That Transparency & Control moments could include screen-based summaries means that not only do users have access to a source of clarity, but also to a definitive record of the conversation that took place.

In addition, it was considered beneficial to allow users to choose how they receive screen-based information. For some users, an immediate transcript or a pop-up notice may suffice; for others, sending it via email may be an appropriate option if documentation is the primary goal.

"Screen-based is helpful because then [a user can] look at it when she wants and have time to review".

- BERLIN



Voice should remain primary

In addition to the dashboard and the transcript, there were many smaller areas where workshop participants felt that visuals would be useful additions, but very rarely did any expect that a visual UI would have to be primary.

For example, there were strong expectations that voice-first consent should be as easy to withdraw as it is to provide. In a voice-specific context, this means that it should be easy to withdraw consent via a voice command, if technical capabilities allow.

Pushing users to a secondary surface via a link, while providing minimal information (e.g., the screen/app) on how to revoke consent, was considered a poorer experience.

"If you can give consent [via voice], you should be able to revoke the same way."

- BERLIN

"I can't open a link in the middle of the street, my phone will get robbed!"

- BRASILIA



Overall recommendations

- Allow privacy settings to be managed via voice
- Use an on-screen summary or confirmation of the decision made.
- Provide users with a central place to review and adjust decisions made via voice.

EXAMPLE CHAT

CONSENT REVOCATION VIA VOICE

Participants showed a strong preference that a voice-first consent should be as easy to withdraw as it is to provide.

REVIEWING VIA SCREEN

Participants argued that users expect to have a centralized visual “transparency decision dashboard” to review their history of choices and legal documents.

Hey are you still using my streaming app ratings to give me recommendations? 

 Yes, I am. Would you like me to recommend a movie for you to stream?

Not right now. Can you stop using my streaming profile for recommendations? 

 Sure. I've changed your preference to no longer use your streaming profile.
I'll also send a notification to your phone so you can review your preferences at any time. Is there anything else I can help you with?

No, that's all. 

4

Summary

Summary

The insights presented in this report represent an early exploration with policy stakeholders around Transparency & Control for AI-powered voice.

As AI capabilities evolve, so too will the expectations of users and the requirements of policymakers. While we have identified emerging opportunities to support Transparency & Control, many areas remain for further exploration.

This work serves as an invitation to product makers and policy stakeholders to continue testing, refining, and debating these concepts to ensure that the next generation of conversational experiences is built on a foundation of trust and transparency.

Explorations	Considerations	Design patterns
 <p>Opportunities for Transparency and Control within Conversational UX</p> <p>Conversational language provides us with opportunities and challenges</p>	<ul style="list-style-type: none"> • Empowering natural dialogue • Elevating personalized control • Unlocking deeper engagement 	<ul style="list-style-type: none"> • User-driven/detailed control of settings • Conversational enquiry
 <p>Encountering a Transparency & Control moment</p> <p>What is an appropriate user experience for voice-based Transparency & Control moments, and should they be adaptable?</p>	<ul style="list-style-type: none"> • Signifying a Transparency & Control moment for a user • Focusing on important content • Optionality for style 	<ul style="list-style-type: none"> • Style • Tone • Sounds • Formality
 <p>Sequencing information</p> <p>How do we prioritise information in a linear voice-based experience, while maintaining a good user experience?</p>	<ul style="list-style-type: none"> • Don't try to share all information at once • Skipping non-critical information 	<ul style="list-style-type: none"> • Critical/non-critical information • Progressive disclosure • Skipping
 <p>Multimodality</p> <p>Do we envisage Transparency & Control being possible via voice only, or will we also need a screen?</p>	<ul style="list-style-type: none"> • A visual dashboard to review • A transcript for clarity • Voice should remain primary 	<ul style="list-style-type: none"> • Product controls • Screen-based confirmations • Additional screen-based messages • User control centre

