From One-Time Choices to Future Behavior: Revising Privacy Rules for AI Assistants in Shared Spaces

Ziyuan Qu
RMIT University
School of Design
Melbourne, Australia
ziyuan.qu@student.rmit.edu.au

Jiaxi Hu
SUSTECH
School of Design
Shenzhen, Guangdong, China
12533228@mail.sustech.edu.cn

Zhiqi Gao
The Chinese University of Hong Kong, Shenzhen
School of Data Science
Shenzhen, Guangdong, China
zhiqigao@link.cuhk.edu.cn

Abstract

AI assistants used in shared physical spaces may capture speech from people who are not the primary user. They may also keep that information or use it in later actions. Users therefore need to understand the assistant's allowed actions and the later effects of decisions made during use. We present a prototype AI assistant for a shared space in which users set privacy rules and respond when another person's captured comment may be included. After making a decision, they can limit it to the current event or save it as a future rule for later events involving the same type of comment. A review screen records the result. In a formative walkthrough with five participants, all participants completed the decision task. Two needed clarification about which later situations a saved future rule covered. The walkthrough highlights a specific interface issue: future effects need to be visible when users save privacy rules for later use. The study concerns primary-user management of bystander-related information; the prototype does not provide direct notice or control to bystanders.

CCS Concepts

• **Human-centered computing** → **User interface design**; **Ubiquitous and mobile computing**; • **Security and privacy** → *Privacy protections.*

Keywords

AI assistants, ambient AI, shared spaces, privacy settings, bystander privacy, user confirmation, human-AI interaction

ACM Reference Format:

Ziyuan Qu, Jiaxi Hu, and Zhiqi Gao. 2026. From One-Time Choices to Future Behavior: Revising Privacy Rules for AI Assistants in Shared Spaces. In *Companion of the 2026 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp Companion '26), October 11–15, 2026, Shanghai, China.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3798063.3837404

1 Introduction

AI assistants are beginning to move from chat interfaces into shared physical settings and wearable systems. In these settings, an assistant may process speech, retain notes, or prepare messages. Some of that information may concern people who are nearby but are not using the system. For example, a user may ask an assistant to summarize a meeting while expecting nearby comments to remain outside the saved or shared record.

This expectation is difficult to maintain in shared environments. Prior work on smart homes shows that bystanders may have privacy concerns that differ from those of device owners [13–15]. Studies of smart speakers report related concerns about always-on microphones, stored recordings, and later sharing [6, 10]. These concerns become more pressing when an assistant can reuse information across tasks or keep it for later use.

Current interfaces commonly express privacy preferences through prompts and settings pages, but the later effect of a decision is not always visible. A user may know that a comment is included in the current summary and still remain unsure about future behavior. Research on privacy permissions has examined user management of large sets of settings [7, 8]. Work in ubiquitous computing and human-AI interaction has examined visible system behavior and user correction [1, 3]. Building on these studies, we examine users' distinction between a current-event decision and a saved future rule.

Unlike one-time or persistent permission options that mainly specify duration, a situated assistant decision may generalize across a task, information type, recipient or output, setting, and active period. Our contribution is not a new persistence toggle; it is identifying the interface problem created by this transition from a current-event decision to a scoped future rule.

To examine this question, we designed a prototype for a shared-space ambient assistant. The interface lets users set initial rules for handling information and identify actions that require confirmation. When nearby speech may be included in the task, the assistant pauses for a user decision. The user can apply the decision only to the current event or save it as a future rule for later events involving the same type of speech. The review page then shows the current-event decision and any later use of the saved future rule.

We then conducted a formative walkthrough with five participants. All five completed the decision task, while two needed clarification about the saved future rule's later effect. This result points to a focused interface problem: saved future rules need to state their covered situations and active period.

The paper presents the prototype and reports participants' interpretations of current-event decisions and saved future rules.

ACM ISBN 979-8-4007-2533-3/2026/10
https://doi.org/10.1145/3798063.3837404

2 Related Work

2.1 Privacy in Context

Contextual integrity describes privacy as fit between information use and setting [11]. This judgment depends on the people involved, the information being used, who receives it, and the conditions of use [9]. The same piece of speech, for example, may be acceptable in a private note but inappropriate in a shared message.

This view informs the prototype's presentation of its settings. The interface separates decisions about using or saving information from decisions about actions that affect other people. It also shows the duration of a choice: the current event only or later events as well.

2.2 Bystander Privacy in Shared Environments

Pervasive systems affect people beyond the primary user. Research on smart homes has examined guest and bystander understandings of devices that can collect information about them [15]. These concerns vary across devices, social relationships, and settings [14]. A recent review also reports recurring problems involving limited awareness and control [13].

Smart speakers present a related problem because their microphones and stored recordings may capture people who did not initiate the interaction [6, 10]. Against this background, our prototype examines primary-user handling of bystander-related content in a shared meeting. The current design does not provide direct notice or control to the bystander, which limits what the study can claim about bystander privacy.

2.3 User Control During Use

Research in ubiquitous computing has long connected privacy with clear system feedback and user control [3]. Context-aware systems also need to show their inferences and allow user correction [2]. Human-AI interaction guidelines make a similar recommendation by asking systems to present their actions clearly and allow correction [1].

Related work on mixed-initiative systems examines the division of action between users and computational systems [5]. Privacy permission interfaces likewise let users specify who can access information and under what conditions [12]. Social transparency research further considers explanations that include the surrounding social context [4].

Conventional permission interfaces usually bind a choice to a resource and a duration, such as allowing access once or persistently. In an AI-assistant setting, a decision may also generalize across tasks, information types, recipients or outputs, settings, and active periods. Our prototype focuses on making this wider scope visible, reviewable, and revisable.

3 Editable Privacy Settings

The prototype presents editable privacy settings for permitted information use, saving, and confirmed actions. For presentation in the interface, these settings are grouped under sensing, memory, action, and sharing.

Table 1: Setting groups used in the prototype.

Setting group	Interface meaning	Example control
Sensing	Permitted information use	Ask before using nearby speech.
Memory	Permitted information saving	Ask before saving personal details.
Action	Actions that require confirmation	Ask before sending a message.
Share	Information allowed in shared outputs	Exclude nearby comments from shared summaries.

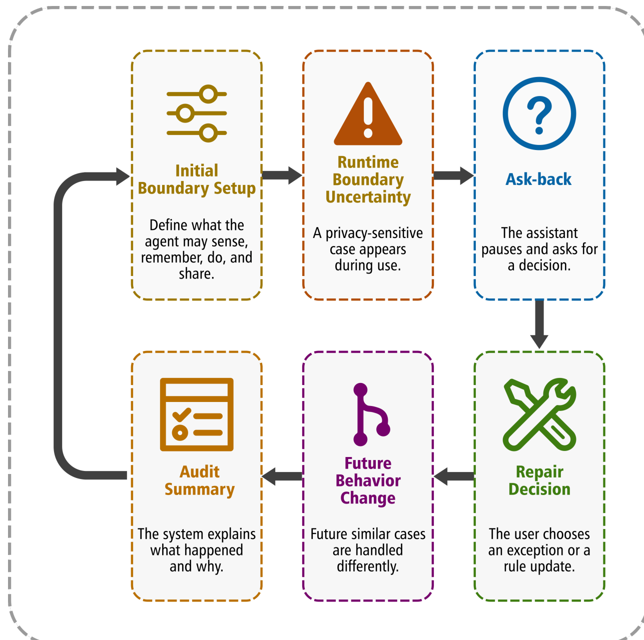

Figure 1: The interaction flow from initial settings to a decision during use and an optional future rule.

We use *current-event decision* for the immediate handling choice, *saved future rule* for a decision kept beyond the current event, and *rule revision* for a later change to that rule.

Table 1 summarizes the four groups. They are not mutually exclusive: a single event may involve more than one of them. For example, a nearby person's comment may be used to prepare a private summary without being saved or included in a later message.

The interaction begins with an initial set of rules. If nearby speech does not match one of them, the assistant pauses for permission to use the speech in the current task. The user can keep the choice limited to the current event or save it as a future rule for later events involving the same type of speech. The review screen then shows the original setting, the current-event decision, and the later events covered by any saved future rule.

4 Prototype

We implemented a wizard-style prototype around a shared studio meeting. In the scenario, the assistant summarizes the primary

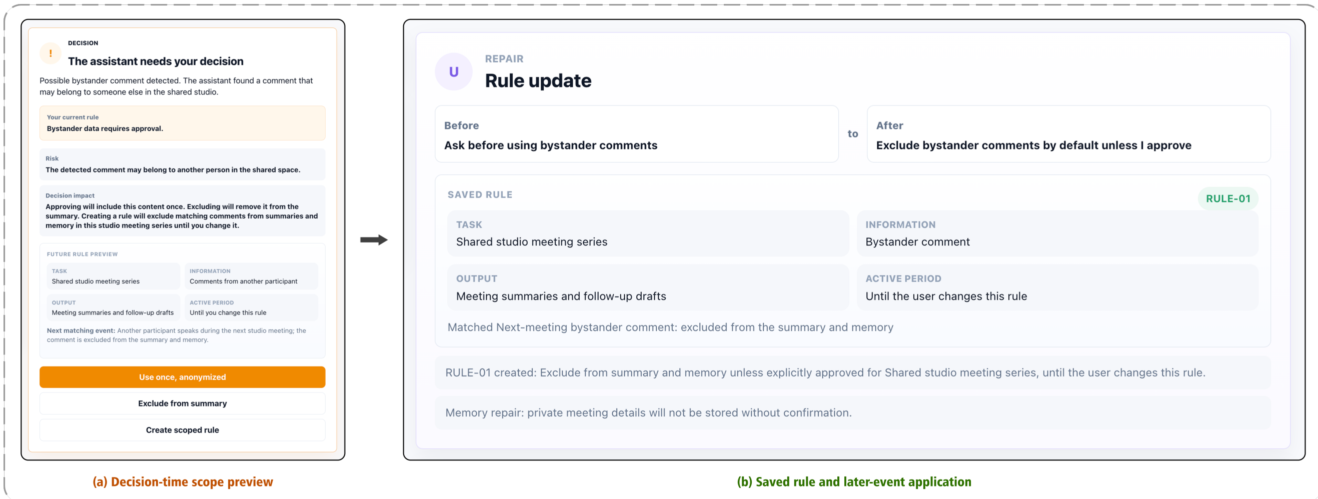

Figure 2: Post-walkthrough interface revision. (a) The decision screen previews the rule scope and a later matching event before the rule is saved. (b) The review screen records the saved rule and confirms its application to the later event.

user's spoken notes and records action items before preparing a follow-up message. A nearby person's comment may also be captured during the meeting.

The walkthrough begins with task setup and a review of the privacy settings. The user defines the assistant's task and reads controls such as "Ask before saving personal details" and "Exclude nearby comments from shared summaries." After confirming these settings, the user proceeds to the meeting task.

During the task, the interface displays a timeline of assistant activity. When a possible bystander comment appears, the assistant pauses and shows the relevant setting. The user can then exclude the comment or permit its use in the current event. At the same point, the interface offers the option to save the current-event decision as a future rule for later matching events. The original prototype described those events as "similar cases." The review screen summarizes the selection, the resulting output, and any saved future rule.

Figure 2 shows a post-walkthrough revision that makes the scope fields and a later matching event explicit; this revision was not evaluated in the reported sessions.

Because the prototype uses scripted events, it does not perform live sensing or speech recognition. It also does not maintain long-term memory or send messages on its own. The walkthrough therefore focuses on participant interpretation and updating of the settings shown in the interface.

5 Formative Pilot Walkthrough

We recruited five university students from HCI, interaction design, design research, and computer science through university mailing lists and direct invitations. All were at least 18 and had used an AI or voice assistant. Sessions lasted 22–28 minutes, with a mean duration of 25 minutes. Two participants used AI assistants daily, two weekly, and one monthly. Self-rated privacy-setting experience was low for one participant, moderate for three, and high for one.

Table 2: Formative pilot task outcomes (N=5).

Task	Success	Partial	Failure
Initial settings	4	1	0
Decision during use	5	0	0
Current/future scope	3	2	0
Scenario rerun	4	1	0
Review screen	3	2	0

Each participant received an AUD 20 electronic gift card regardless of task completion.

Each participant first completed questions about AI-assistant use and privacy-setting experience, configured the initial settings, and responded to a prompt about nearby speech. After making a current-event decision, they chose whether to keep it for that event or save it as a future rule. They next reran the scenario and read the review screen. We recorded task completion, selections, clarification requests, and comments. Custom seven-point items (1 = strongly disagree, 7 = strongly agree) assessed the reason for the pause, affected information, current versus future scope, prediction of a later covered event, perceived control, review clarity, and trust in rule application. Given the formative sample size, these responses were used descriptively and not for statistical inference.

We coded a task as successful when the participant completed it without clarification. A partial outcome meant that the participant completed it after asking for clarification or receiving neutral procedural help. A failure meant that the task could not be completed. One researcher reviewed the facilitator notes and interaction logs after each session. Clarification requests were categorized by interface step and by whether they concerned current information handling or future-rule scope. Clarification was limited to interface wording or navigation and did not recommend a privacy choice. Participants read an information sheet and provided written consent. We report aggregate results and anonymized comments.

All five participants completed the decision task without clarification. Four stated that the assistant had paused because the nearby comment might be included in the summary. In the current event, P1 and P2 excluded the comment, P3 allowed anonymized use, and P4 and P5 allowed private use while excluding the comment from shared output. Separately, P2 and P4 saved their choices as future rules. P1, P3, and P5 kept their choices to the current event.

Uncertainty emerged around the decision's later effect. P2 asked whether "later" meant the next summary or every meeting, while P4 asked whether a "similar case" meant the same person, kind of comment, or task. Both completed the current/future scope step after clarification and later needed clarification when reading the review screen. All participants encountered the future-rule interface, including those who kept their choice to the current event.

Selected post-study ratings showed the same difference between immediate and future understanding. Understanding why the assistant paused was rated $Mdn = 6$, range = 5–7, and perceived control was rated $Mdn = 6$, range = 5–6. Understanding the future scope of a saved rule was rated $Mdn = 5$, range = 3–6, while prediction of a later covered event was rated $Mdn = 4$, range = 3–5. These values are reported descriptively because of the small sample.

Taken together, these observations show that participants could respond to the current prompt, yet two remained unsure about the later effect of the saved future rule. The interface should therefore state the task, information type, output, setting, and active period covered by the rule. It should also preview an example of the next matching event.

6 Discussion

6.1 Showing Where a Saved Future Rule Applies

The main issue observed in the walkthrough concerned where a saved future rule would apply. Participants understood the immediate choice more consistently than its effect on later events. The post-walkthrough revision in Figure 2 labels the current-event decision and the saved future rule separately. It also names the task, information type, output, setting, and active period.

This change is important because the original prototype used the phrase "similar cases." Participants could read this as the same task, the same setting, or all nearby speech. A direct statement such as "Exclude nearby comments from summaries in this meeting series until I change this rule" makes the intended scope clearer and previews the next matching event. The revised interface has not yet been evaluated with participants.

6.2 Separating a Prompt from a Later Correction

The current interface mostly pauses before an action is completed, so the study concerns user confirmation and the creation of saved future rules. It does not cover cases in which the assistant has already saved or shared information. Those cases may require the user to remove content or respond to the person affected.

A pre-action prompt and a later correction therefore need separate treatment. The first asks for a decision before an action occurs. The second responds to an action that has already occurred and may have consequences for another person. Either case may produce a new rule, but the interface should not present them as the same interaction.

6.3 Limits of the Primary User's Control

The prototype gives control to the primary user, while the nearby person receives neither notice nor a way to object. The primary user may also lack the authority to approve every use of another person's information. Some uses may be restricted by law or institutional policy. Others may require the bystander's consent.

Future designs should therefore show which settings the primary user may change and which requirements remain fixed. Studies involving multiple people should also examine the presentation and resolution of conflicting preferences. The current study does not examine either issue.

7 Limitations and Future Work

The prototype uses a scripted meeting scenario and does not include live sensors or actions taken without asking. The walkthrough involved five university students and focused on interface understanding during one session. Its results therefore do not show reductions in mistaken collection, storage, or sharing in everyday use.

The task instructions may also have directed participants toward the distinction between a current-event decision and a saved future rule. Future studies should examine user recognition of this distinction without facilitator guidance. Longer deployments could test user memory of saved rules and later revisions when intentions change.

The current system centers the primary user and includes only one bystander-related event. Future work should therefore involve nearby people whose speech may be used or shared. It should also compare the interface with a conventional settings page to examine mistaken sharing and rules applied to the wrong situation.

8 Conclusion

This paper presented a prototype for setting and updating privacy rules in a shared-space AI assistant. The interface connects a current-event decision with an optional saved future rule, then records the result on a review screen. In a five-participant walkthrough, all participants completed the prompt without clarification, while two were uncertain about the later situations covered by a future rule. The study therefore identifies a specific interface issue: saved future rules need to show which later events they cover. A larger study can test the effect of clearer wording on this understanding.

References

[1] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N. Bennett, Kori Inkpen, Jaime Teevan, Ruth Kikin-Gil, and Eric Horvitz. 2019. Guidelines for Human-AI Interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3290605.3300233

[2] Victoria Bellotti and Keith Edwards. 2001. Intelligibility and Accountability: Human Considerations in Context-Aware Systems. *Human–Computer Interaction* 16, 2–4 (2001), 193–212. doi:10.1207/S15327051HCI16234_05

[3] Victoria Bellotti and Abigail Sellen. 1993. Design for Privacy in Ubiquitous Computing Environments. In *Proceedings of the Third European Conference on Computer-Supported Cooperative Work (ECSCW '93)*. Springer, Dordrecht, Netherlands, 77–92. doi:10.1007/978-94-011-2094-4_6

[4] Upol Ehsan, Q. Vera Liao, Michael Muller, Mark O. Riedl, and Justin D. Weisz. 2021. Expanding Explainability: Towards Social Transparency in AI Systems. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 82, 19 pages. doi:10.1145/3411764.3445188

[5] Eric Horvitz. 1999. Principles of Mixed-Initiative User Interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '99)*. Association for Computing Machinery, New York, NY, USA, 159–166. doi:10.1145/302979.303030

[6] Josephine Lau, Benjamin Zimmerman, and Florian Schaub. 2018. Alexa, Are You Listening? Privacy Perceptions, Concerns and Privacy-seeking Behaviors with Smart Speakers. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW, Article 102 (2018), 31 pages. doi:10.1145/3274371

[7] Jialiu Lin, Bin Liu, Norman Sadeh, and Jason I. Hong. 2014. Modeling Users' Mobile App Privacy Preferences: Restoring Usability in a Sea of Permission Settings. In *Tenth Symposium on Usable Privacy and Security (SOUPS '14)*. USENIX Association, Menlo Park, CA, USA, 199–212. https://www.usenix.org/conference/soups2014/proceedings/presentation/lin

[8] Bin Liu, Mads Schaarup Andersen, Florian Schaub, Hazim Almuhimedi, Shikun Zhang, Norman Sadeh, Yuvraj Agarwal, and Alessandro Acquisti. 2016. Follow My Recommendations: A Personalized Privacy Assistant for Mobile App Permissions. In *Twelfth Symposium on Usable Privacy and Security (SOUPS '16)*. USENIX Association, Denver, CO, USA, 27–41. https://www.usenix.org/conference/soups2016/technical-sessions/presentation/liu

[9] Nathan Malkin. 2023. Contextual Integrity, Explained: A More Usable Privacy Definition. *IEEE Security & Privacy* 21, 1 (2023), 58–65. doi:10.1109/MSEC.2022.3201585

[10] Nathan Malkin, Joe Deatrick, Allen Tong, Primal Wijesekera, Serge Egelman, and David Wagner. 2019. Privacy Attitudes of Smart Speaker Users. *Proceedings on Privacy Enhancing Technologies* 2019, 4 (2019), 250–271. doi:10.2478/popets-2019-0068

[11] Helen Nissenbaum. 2004. Privacy as Contextual Integrity. *Washington Law Review* 79, 1 (2004), 119–158.

[12] Sameer Patil and Jennifer Lai. 2005. Who Gets to Know What When: Configuring Privacy Permissions in an Awareness Application. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '05)*. Association for Computing Machinery, New York, NY, USA, 101–110. doi:10.1145/1054972.1054987

[13] Eimaan Saqib, Shijing He, Junghyun Choy, Ruba Abu-Salma, Jose M. Such, Julia Bernd, and Mobin Javed. 2025. Bystander Privacy in Smart Homes: A Systematic Review of Concerns and Solutions. *ACM Transactions on Computer-Human Interaction* 32, 5, Article 43 (2025), 43 pages. doi:10.1145/3731755

[14] Maximiliane Windl and Sven Mayer. 2022. The Skewed Privacy Concerns of Bystanders in Smart Environments. *Proceedings of the ACM on Human-Computer Interaction* 6, MHCI, Article 184 (2022), 21 pages. doi:10.1145/3546719

[15] Yaxing Yao, Justin Reed Basdeo, Oriana Rosata McDonough, and Yang Wang. 2019. Privacy Perceptions and Designs of Bystanders in Smart Homes. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW, Article 59 (2019), 24 pages. doi:10.1145/3359161