Speculative Design Opportunities in AI-Mediated Privacy Management

LISA MEKIOUSSA MALKI

ACM Reference Format:
Lisa Mekioussa Malki. 2026. Speculative Design Opportunities in AI-Mediated Privacy Management. In *Proceedings of The ACM (Association of Computing Machinery) CHI conference on Human Factors in Computing Systems (CHI '26).* ACM, New York, NY, USA, 3 pages. https://doi.org/XXXXXXX.XXXXXXX

1 Introduction

The paradigm of notice and choice has dominated online privacy regulation for several decades. Under this framework, data processors must provide transparency notices and consent mechanisms intended to enable informed decision-making about personal data use [1]. However, informed consent is undermined by socio-technical power imbalances in the digital sphere: transparency notices are often lengthy, opaque, and filled with technical jargon, while choice mechanisms are frequently poorly designed and difficult to exercise meaningfully. These limitations suggest that notice and choice is insufficient for supporting user autonomy, motivating the need for new regulatory and design paradigms. While prior work has explored the automation of privacy management, recent advances in agentic large language models (LLMs) enable more sophisticated forms of support and personalisation. Personalised Privacy Assistants (PPAs) are conceptualised as intelligent agents that learn user privacy preferences and configure settings automatically or semi-automatically [3]. By shifting privacy management away from frequent, cognitively demanding decisions, LLM-enabled PPAs may reduce the burden of managing complex settings [8]. However, AI assistants that continuously observe and act on user behaviour are inherently surveillant, and LLMs may engage in manipulation or coercion to encourage users to disclose sensitive information [9]. This paper presents findings from a speculative design exercise illustrating how these tensions may manifest in the future, and contributes a user persona that foregrounds privacy, power, and human autonomy in AI-mediated privacy management.

2 Design Provocation: A Browser-Based Privacy Assistant

To conduct the analysis, I performed a speculative brainstorming exercise guided by the workshop prompt : *"How can emerging interaction modalities such as voice interfaces, generative AI, or social robots reshape how people engage with consent and control?"* The exercise was organised around three research questions:

(1) What are the existing challenge areas associated with notice and choice?
(2) How might current and emerging AI agent capabilities address these problems?
(3) What consequences and risks emerge from applying AI in this context?

Author's Contact Information: Lisa Mekioussa Malki.

Manuscript submitted to ACM

I followed a multi-step approach, beginning with a review of relevant literature, articles, and existing technical tools. Key insights were synthesised into themes reflecting major challenges in notice and choice, and how capabilities of agentic AI could be integrated into a browser-based privacy assistant. I then developed a future-oriented narrative persona to illustrate how a user might interact with such a system and how concerns around autonomy and trust may arise. The persona draws on the Futures Design Toolkit, and aggregates themes identified in the literature into an experiential 'day-in-the-life' narrative. The full set of findings is documented in a Miro board[1].

Opportunity 1: Dynamising Privacy Communications. The shortcomings of privacy policies are well documented: they are lengthy, complex, and difficult for the public to understand. In the near future, LLMs may summarise these documents and extract key practices and data rights, sparing users the need to navigate pages of dense text. Conversational agents could transform notice from static disclosure into ongoing, scaffolded privacy sensemaking. Instead of passively absorbing information, users may instead engage in multimodal dialogue with an AI agent to ask questions, explore implications, and obtain personalised explanations [4].

Opportunity 2: Automating Browser Navigation for User-Friendly Privacy Control. The contemporary digital landscape requires individuals to interact with dozens of websites, making privacy management infeasible at scale. Multimodal LLMs offer the potential to autonomously navigate browser GUIs and handle consent tasks by:

(1) Interpreting cookie banners, consent dialogs, and privacy policies [4].
(2) Translating high-level preferences into concrete decisions and reasoning qualitatively across legal/social norms [5].
(3) Using Document Object Model (DOM) and visual understanding to identify and activate controls across diverse UIs, including those employing deceptive patterns.
(4) Update or revoke consent in response to policy changes or evolving preferences.

Once a privacy choice is identified, different interaction models are possible: the agent may provide recommendations, act autonomously, or take a hybrid approach [7]. For example, it could restructure consent workflows, condensing long vendor lists into meaningful categories presented via a conversational or dynamically generated UI.

Opportunity 3: Cross-Site Privacy Actions. The frequent evolution of data practices and user privacy needs necessitates dynamic, ongoing consent [2]. A speculative privacy agent might continuously monitor the broader privacy landscape by accessing up-to-date information feeds (e.g., via retrieval-augmented generation), scanning for changes to privacy policies or reports of problematic practices, and distilling this information in order to help steer users toward more privacy-respecting services. This may be realised through global digital hygiene actions, such as a deletion 'kill switch' that removes all user data from a high-risk service via interface automation, as well as more granular commands such as *"Close accounts I haven't used in two years"* or *"Delete all social media posts relating to my children"* [8].

Potential Challenges: Over-reliance, Surveillance, and Manipulation. Although these capabilities may reduce harmful data practices and improve regulatory enforcement, delegating privacy decisions raises concerns [7]. As internet use becomes increasingly mediated by AI, users may become alienated from their own digital lives, with privacy management shifting into another black box. Furthermore, to act effectively, agents require detailed insight into users' preferences and behaviours. How this information is obtained introduces further tension: agents may infer preferences from behavioural data such as past decisions, or they may directly elicit attitudinal information from users [3]. However,

[1] https://miro.com/app/board/uXjVGIgH3hg=/?share_link_id=14236500442

privacy preferences themselves can be sensitive, revealing intimate aspects of identity and vulnerability [6]. Possible mitigations include the minimal and tightly controlled use of passively acquired data, and local processing to ensure that sensitive information does not leave the user's device. Prior user studies on AI-enabled PPAs consistently suggest that such tools should assist rather than replace human decision-making, and must provide clear mechanisms for transparency and oversight [3, 8] An open-source, value-aligned development model may also help prevent agents from steering users toward invasive choices [7]. The following user persona illustrates many of these tensions.

> In a world where AI agents handle browser-based privacy decisions, Effie barely has to think about privacy anymore. She started using the PrivacyFirst browser last month: after onboarding through a voice-based 'privacy interview', it now manages almost everything for her. From automatically dealing with cookie banners to shutting down entire accounts with a single command, privacy just sort of happens in the background. For example, she recently read an article saying that a site she uses will start selling user data to AI companies. This made her really anxious, so she asked her agent to find a more privacy-friendly alternative and delete her old account after exporting all her data. Still, Effie isn't completely at ease. Sometimes it feels like the agent knows her better than she knows herself, and she'd be pretty lost without it. She also worries that the agent will make a mistake like deleting the wrong account, but she's reassured by the fact that it always checks with her before making big decisions. Since it's maintained by a non-profit group of regulators and expert developers she generally trusts it.

3 Conclusion and Further Discussion

In conclusion, while agentic privacy assistants can ease cognitive burden and help counter power imbalances in notice-and-choice frameworks, they may also introduce risks of surveillance, manipulation, and user dependency. Through speculative design, this paper highlights the tensions that emerge when privacy decision-making is delegated to AI. Several questions arise for further discussion: How can AI-mediated privacy systems support users without undermining their agency? How can agents use to learn preferences without becoming intrusive or extractive? What governance, technical, or institutional safeguards (e.g., local processing) are necessary to ensure PPAs are trustworthy?

References

[1] Lorrie Faith Cranor. 2024. Notice and Choice Cannot Stand Alone. *Commun. ACM* 67, 12 (Nov. 2024), 37–39. doi:10.1145/3699527

[2] Lisa Mekioussa Malki, Ina Kaleva, Dilisha Patel, Mark Warner, and Ruba Abu-Salma. 2024. Exploring Privacy Practices of Female mHealth Apps in a Post-Roe World. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 576, 24 pages. doi:10.1145/3613904.3642521

[3] Victor Morel, Leonardo Horn Iwaya, and Simone Fischer-Hübner. 2025. AI-Driven Personalized Privacy Assistants: A Systematic Literature Review. *IEEE Access* 13 (2025), 160982–161002. doi:10.1109/ACCESS.2025.3609188

[4] Bolun Sun, Yifan Zhou, and Haiyun Jiang. 2024. Empowering Users in Digital Privacy Management through Interactive LLM-Based Agents. doi:10.48550/arXiv.2410.11906

[5] Brian Wang, Luis Antonio Garcia, and Mani Srivastava. 2024. PrivacyOracle: Configuring Sensor Privacy Firewalls with Large Language Models in Smart Built Environments. In *2024 IEEE Security and Privacy Workshops (SPW)*. doi:10.1109/SPW63631.2024.00028

[6] Wen Wang and Beibei Li. 2024. Learning Personalized Privacy Preference from Public Data. *Information Systems Research* 36 (06 2024). doi:10.1287/isre.2023.0318

[7] Meihe Xu, Arianna Rossi, and Aurelia Tamò-Larrieux. 2025. The Future of Personalized Privacy Assistants: Gathering of Expert Opinions. *Digital Society* 4, 3 (Dec. 2025), 75. doi:10.1007/s44206-025-00232-4

[8] Meihe Xu, Aurelia Tamò-Larrieux, and Arianna Rossi. 2025. Acceptability of AI Assistants for Privacy: Perceptions of Experts and Users on Personalized Privacy Assistants. doi:10.48550/ARXIV.2509.08554 Version Number: 1.

[9] Xiao Zhan, Juan Carlos Carrillo, William Seymour, and Jose Such. 2025. Malicious LLM-Based Conversational AI Makes Users Reveal Personal Information. doi:10.48550/arXiv.2506.11680 arXiv:2506.11680 [cs].