Defining the Foundations for Consent with AI Systems

PATRICK GAGE KELLEY, Google, USA

While there may come a time where are closer to actual deployments of complex rules for agentic privacy decision making or novel AI consent modalities – I ponder in this short paper what the foundational values to allow people to give free, informed consent for AI systems should be. That is, can we even agree upon the set of basic, simple affordances that every person interacting with an AI system should have?

Additional Key Words and Phrases: consent, artificial intelligence, AI, LLM, transparency, control

ACM Reference Format:
Patrick Gage Kelley. 2026. Defining the Foundations for Consent with AI Systems. In *Proceedings of CHI 2026 (CHI '2026).* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

In 2023 my team published a paper [1] that analyzed 9,765 open-ended responses to the question:

> 'Now we would like to ask you to think about Artificial Intelligence (AI) and privacy. In what ways will Artificial Intelligence (AI) affect privacy in the future? Please be specific.'

We coded those responses using a codebook of 368 codes, 36 of which were focused on privacy topics, and found four buckets of privacy concerns. Those were: **Highly Personal, Data at Risk, State & Surveillance,** and **Without Consent.** Without consent, while the least frequent of the four, was quite interesting to me. We found our respondents were quite accurately able to describe the reasons why AI was already beginning to push beyond their comfort zone, with respect to privacy (then a more abstracted concept, as ChatGPT had not yet become public, not yet a part of everyday life).

We described three insights from our section on consent, which I briefly summarize again here.

1. Data gathering and use occur without consent. – our respondents felt that true consent to have their data kept from AI systems, or to limit the use of AI systems to make decisions or discover things about them, was impossible:

> People's data will be invaded whether they know and give their consent or not. –Kenya

> Invasion of privacy by spying on me without my consent –South Korea

Author's Contact Information: Patrick Gage Kelley, patrickgage@acm.org, Google, New York City, New York, USA.

Manuscript submitted to ACM

2. Unaware. – not only would people be surveilled by AI without realizing it, but respondents recognized that it would become more and more difficult to understand how AI systems operated, collected and processed information, and as a result lead to people being unaware of what possible control or agency they could have:

> It has already invaded households beyond what the majority of people know. There is no privacy now. –United States

> We will not have privacy anymore. Companies will use our data and we won't even know. –Brazil

3. Personal information required. – our respondents felt that without giving up their personal information, they wouldn't be able to use AI – that is, they wouldn't be truly giving consent freely, but rather be forced to do so:

> Useful features will be available in exchange for the disclosure of personal information. –Japan

> Artificial intelligence requires people to reveal themselves, while inevitably exposing their privacy –China

While these responses were collected before Generative AI became commonplace, in many ways they feel even more urgent today. Many of the concerns they describe above still apply, and few solutions have been deployed to assist users exercise a more true form of consent.

Also back in 2023, my team published an online rubric to improve *explainability* [2], which can be viewed online at https://explainability.withgoogle.com/rubric. This rubric includes many items that help improve user's AI literacy, and understanding for particular systems that implement them, but also provide what I think of as the most meager baseline for helping users have more agency over consent. This includes items such as **On/Off**: "Describe when the AI feature is on/off, whether the user can turn the feature on/off, and whether/how the user has visibility of an AI feature running in the background." But even giving users the ability to (easily) turn an AI system off is not standard today.

As we are "Rethinking Consent and User Control in the Age of AI" I would like to bring a conversation of underlying values and baseline standards to the conversation at this workshop. Building off the Explainability Rubric and our public opinion polling surveys of AI privacy concerns, and in conversation with the experts at the workshop, I believe we can make good headway towards an initial baseline for consent and user control that AI systems should support.

Short Bio

Patrick Gage Kelley is a Research Lead for AI Research & Standards on Google's Trust & Safety team. He has worked on projects that help us better understand how people think about their data and safety online. These include projects on the use and design of user-friendly privacy displays, passwords, location-sharing, mobile apps, encryption, technology ethics, designing products for people with the most significant digital safety risks, and most recently on people's relationship and understanding of AI.

References

[1] Patrick Gage Kelley, Celestina Cornejo, Lisa Hayes, Ellie Shuo Jin, Aaron Sedley, Kurt Thomas, Yongwei Yang, and Allison Woodruff. 2023. "There will be less privacy, of course": how and why people in 10 countries expect AI will affect privacy in the future. In *Proceedings of the Nineteenth USENIX Conference on Usable Privacy and Security* (Anaheim, CA, USA) *(SOUPS '23)*. USENIX Association, USA, Article 32, 25 pages.

[2] Patrick Gage Kelley and Allison Woodruff. 2023. Advancing Explainability Through AI Literacy and Design Resources. *Interactions* 30, 5 (Aug. 2023), 34–38. doi:10.1145/3613249