

Relational Consent Scaffolding: Calibrated Reliance and Psychological Ownership in AI-Driven User Control

Position paper for CHI '26 Workshop: Moving Beyond Clicks — Rethinking Consent and User Control in the Age of AI

DAVID SARLOS, Dyson School of Design Engineering, Imperial College London, United Kingdom

Many digital consent experiences reduce complex data practices to a single click, offering little agency and obscuring the role of AI. Building on Psychological Ownership theory and research on calibrated reliance, this paper proposes *relational consent scaffolding*: a design framework that pairs reflective modalities with explainable AI assistance to support informed, contextual and delegated consent. Instead of a transactional click, consent becomes a process of understanding, accountability and adjustable delegation. The framework outlines design principles, explores when AI may assist or automate decisions and sketches consequences for users and service providers, contributing to the workshop's agenda of moving beyond clicks.

1 Introduction

Many digital consent experiences reduce complex data practices to a single click, leaving people uninformed and unable to control how AI systems use their information. Psychological Ownership (PO) research emphasises that feelings of “mine” arise through control, self-investment and knowledge of a resource [4], while misaligned affordances can diminish users' sense of ownership over their data. These insights, combined with evidence that richer engagement formats elicit deeper care [5], motivate a move away from scroll-and-click consent.

This paper proposes *relational consent scaffolding*, a framework that treats consent as a relationship rather than a transaction. Building on PO design insights from circular possessions [2], this framework suggests that explainable AI support can help users progressively build ownership over both their data and their consent choices. In line with the workshop's agenda, the approach considers when AI may act on a user's behalf, how delegation can reduce cognitive burden without undermining autonomy, and what safeguards distinguish assistance from manipulation.

2 Theoretical framing and background

2.1 Psychological Ownership and digital consent

PO research posits that individuals feel a resource is “theirs” when they can exercise control over it, invest themselves in it and develop intimate knowledge of it [4]. Conversely, when digital affordances are ambiguous or misaligned, possessions can feel contested. Translating these insights to data consent suggests that meaningful consent should offer opportunities to manage data flows, invest attention and understand consequences rather than requiring a single tap. A complementary way to make consent ‘situated’ is to treat it as part of a behaviour setting—shaped by roles, norms, routines, and infrastructure—rather than an isolated individual choice. Setting-Driven Design provides a practical framing for mapping these setting elements and identifying intervention points for changing consent routines without relying on repeated click-through decisions [1].

2.2 Calibrated reliance on AI agents

The PhD research underpinning this workshop paper explores conditions that support *calibrated reliance*: relying on AI assistance when it improves decision quality and resisting it when it does not [6]. Through studies of video-based

Author's Contact Information: David Sarlos, Dyson School of Design Engineering, Imperial College London, London, United Kingdom, david.sarlos@imperial.ac.uk.

interactions and provenance cues, we plan to examine how people interpret uncertain information and decide whether to act. A central premise of this approach is that AI can be leveraged to summarise and contextualise policies, provided it does not supplant human agency. Framed through Setting-Driven Design, relational consent scaffolding leverages AI as an agent playing an advisory role, translating opaque legal policies into contextual summaries and highlighting data practices misaligned with user-defined consent preferences, all while specifying clear thresholds for human-only interaction and control.

3 Relational consent scaffolding

3.1 Design principles

The proposed framework comprises four principles:

- (1) **Reflective modalities.** Consent flows should invite reflection rather than immediate acceptance. Deploying short text narratives or setting-driven video summaries contextualises how data is used and why it matters, fostering considered judgement and self-investment [5].
- (2) **Adaptive AI assistance.** AI agents can summarise policies and highlight salient risks or benefits, with clear explanations and the ability for users to query decisions. Explainability—meaning the agent actively exposes its reasoning and the specific policy data it references—and avoidance of coercive nudging are essential.
- (3) **Delegation with oversight.** Users may delegate routine decisions with clear scope and revocation. Because risk is subjective, the threshold between automated low-stakes actions and high-stakes active confirmation is established by a custom agent operating within user-defined context protocols.
- (4) **Ownership cues and accountability.** Interfaces should geographically visualise data flows and their jurisdictional implications, clarifying the specific rights users can exercise. A consent journal that logs and revisits decisions supports ongoing control and review.

Together, these principles transform consent from a transactional hurdle into a reflective and accountable process. By shifting the burden from in-the-moment compliance to scheduled reviews of historical logs, it ensures that explainable AI delegation remains bounded and sensitive to the user’s setting.

3.2 Answering the workshop’s key questions

When can active privacy decision-making be lifted? Active decision-making may be lifted when the context is low-stakes, the data requested is proportionate to the service’s functionality, and the user has explicitly authorised delegation to verified providers within defined legislative jurisdictions. For example, an EU-based user might configure their AI agent to automatically decline all marketing trackers except those from local, GDPR-compliant providers, freeing them from repetitive choices. Ethically, this relies on clear specification of delegation scope, transparent oversight of anomalous agent decisions, and immediate revocation rights.

To what extent can an AI agent decide on a user’s behalf? AI agents can filter options, translate policies and recommend actions, but final authority should remain with the user in high-impact scenarios—defined as anomalous requests, actions exceeding adaptive thresholds, or the processing of highly sensitive medical and financial data. In the proposed scaffolding, AI agents act as advisors or proxies with bounded agency; they cannot consent to sensitive data transfers or legal agreements without explicit sign-off. This preserves autonomy while distributing cognitive load.

How might delegation reduce burden while preserving agency? By allowing users to configure tiers of delegation (e.g., automated for routine preferences, manual for sensitive domains), the framework reduces repetitive tasks but requires occasional review. This hybrid model respects the user’s right not to be overwhelmed while maintaining oversight. The “consent journal” and periodic check-ins are designed to ensure that delegation does not drift into unchecked automation.

What safeguards distinguish supportive guidance from manipulation? Safeguards include clear AI explanations, audit trails, and alignment with Psychological Ownership cues. For instance, to preserve Control—a core route to psychological ownership—if an agent highlights the benefits of sharing data, it must symmetrically state the potential risks. Interfaces should strictly avoid dark patterns, such as using emotional appeals to coerce consent. When such manipulative behaviours are encountered, they can be flagged by users or AI agents and shared across collaborative networks, transforming isolated feedback loops into community-wide safeguards supported by regular third-party audits.

4 Consequences

4.1 Aspects shaping today’s consent landscape

Current consent practices are shaped by legal requirements (e.g., GDPR), mass adoption of click-through policies, and a commercial ecosystem that exploits the ‘dark side’ of Psychological Ownership [2] through territorial data hoarding. Users often experience choice overload and “consent fatigue,” leading to habituated acceptance. Digital possessions like personal data suffer from affordance misalignments [3]: individuals lack visibility into where data travels and have limited tools to exercise ownership. AI systems compound these challenges by inferring sensitive traits and generating predictive insights without explicit permission.

4.2 Anticipated consequences of relational scaffolding

Relational consent scaffolding aims to transform consent from a legal formality into an ongoing engagement. By providing control, investment and understanding opportunities, the approach may support stronger feelings of ownership and more considered decisions. Reflective modalities could make data flows more tangible, and adaptive agents may reduce cognitive burden while preserving agency. However, while introducing reflective friction demands a non-negligible time investment that may slow interactions in contexts where speed is valued, this burden is ultimately offset by the AI agent automating high-frequency, routine choices. Poorly designed explanations could overwhelm or mislead, and responsibility would shift: designers need to curate supportive narratives, maintain accountability and align delegation with users’ values. Furthermore, to foster trust and prevent adversarial gaming, these agents should be structured around auditable Model Context Protocols (MCPs) and subjected to independent assurance, guaranteeing they serve the individual’s settings rather than platform incentives. A futures-design lens helps anticipate secondary and tertiary consequences of consent interventions (e.g., whose voices are privileged, and how norms shift across services).

5 Biography and interests

David Sarlos is a design engineer researcher at Imperial College London and founder of Vialog. Drawing on his CHI 2026 paper on Psychological Ownership, he designs mechanisms that move humans toward relational engagement with AI using a setting-sensitive agenda.

References

- [1] Weston Baxter, Peter Mandeno, Robert Aunger, and Edward Brial. 2025. Setting-driven design: a context-driven approach to behavioural design. *Design Science* 11 (2025), e35. doi:10.1017/dsj.2025.10025
- [2] Weston L Baxter. 2017. Designing circular possessions: exploring human-object relationships in the circular economy. (March 2017). doi:10.25560/52779
- [3] Rebecca Mardon, Janice Denegri-Knott, and Mike Molesworth. 2023. “Kind of Mine, Kind of Not”: Digital Possessions and Affordance Misalignment. *Journal of Consumer Research* 50, 2 (July 2023), 255–281. doi:10.1093/jcr/ucac057
- [4] Jon L. Pierce, Tatiana Kostova, and Kurt T. Dirks. 2003. The State of Psychological Ownership: Integrating and Extending a Century of Research. *Review of General Psychology* 7, 1 (March 2003), 84–107. doi:10.1037/1089-2680.7.1.84
- [5] David Sarlos and Weston Baxter. 2026. From Participation to Relational Engagement: Psychological Ownership in Digital Petitions. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI'26)*. ACM, Barcelona, Spain, 21 pages. doi:10.1145/3772318.3791561
- [6] Max Schemmer, Niklas Kuehl, Carina Benz, Andrea Bartos, and Gerhard Satzger. 2023. Appropriate Reliance on AI Advice: Conceptualization and the Effect of Explanations. In *Proceedings of the 28th International Conference on Intelligent User Interfaces*. ACM, Sydney NSW Australia, 410–422. doi:10.1145/3581641.3584066