EPH–International Journal of Educational Research
Published by: **Southern Cross Academic Press**
ISSN (Online): **2208-2204**
Volume 10 Issue 01 April 2026

CONSENT FATIGUE IN CONVERSATIONAL AI: MICRO-DESIGN REMEDIES FOR OVEREXPOSED USER

Shweta Chaturvedi[1*]
[1]Master of Laws Dhirubhai Ambani University School of Law, Email Id:Shweta@Solacelawpractice.com
Orcid Id: 0009-0004-6714-6132

Article History:

Article Type:
Received Date:**26//02/2026**
Revised Date: **31//03/2026**
Accepted Date:**08//04/2026**
Published Date: **15/04/2026**

Keywords: Consent Fatigue; AI systems; human-computer interaction

ABSTRACT

The operationalisation of informed consent in conversational AI systems has produced a paradox: the more assiduously platforms comply with notification obligations, the less meaningful individual consent becomes. Repeated exposure to privacy prompts, disclosure notices, and permission requests within chatbot and voice assistant interfaces engenders a condition this article terms 'consent fatigue'—a cognitive and behavioural disposition characterised by habituated disengagement from the very disclosures designed to protect user autonomy. Drawing upon a mixed-methods study involving one hundred and twenty-seven participants across varied demographic groups, this article investigates the relative efficacy of micro-design interventions—including icon-based cues, adaptive frequency controls, and contextual micro-prompts—in ameliorating consent fatigue without sacrificing regulatory compliance. The findings demonstrate a statistically significant inverse correlation between traditional text-heavy consent mechanisms and user comprehension, and a correspondingly significant improvement in both engagement and understanding under icon-based and adaptive designs. The article argues that current regulatory frameworks under the GDPR and the California Consumer Privacy Act (CCPA) are inadequate to the human-computer interaction realities of conversational AI, and proposes a revised consent design standard grounded in cognitive load theory, behavioural economics, and user-centred legal design.

1. Introduction

Informed consent occupies a foundational position in contemporary data protection law.[Regulation (EU) 2016/679 (GDPR), Art 7(2); Recital 32.] The principle that individuals should exercise meaningful control over the collection and processing of their personal data underlies the consent provisions of the GDPR, the CCPA, and cognate instruments across common law and civil law jurisdictions alike.[California Consumer Privacy Act of 2018, Cal Civ Code ss 1798.100-1798.199 (CCPA), s 1798.120.] Yet the practical realisation of this principle in the context of conversational AI systems—chatbots, virtual assistants, and voice-activated platforms—has proved far more elusive than legislative drafters anticipated.

The difficulty is not primarily technical but cognitive. Privacy notices and consent prompts are, by their nature, interruptions. In the context of a conversational AI interaction—where users engage in continuous, real-time dialogue—interruptions for consent purposes are particularly disruptive. Studies in human-computer interaction have long demonstrated that repeated, uncontextualised interruptions degrade task performance and increase the probability of habituated non-response.[Aleecia McDonald & Lorrie Faith Cranor, 'The Cost of Reading Privacy Policies' (2008) 4 I/S: A Journal of Law and Policy for the Information Society 543, 561.] Applied to consent, this dynamic produces what scholars have termed 'notice fatigue' or, in the specific register of AI interaction, 'consent fatigue': the progressive disengagement of users from consent mechanisms through sheer exposure.[Lorrie Faith Cranor, 'Necessary but Not Sufficient: Standardized Mechanisms for Privacy Notice and Choice' (2012) 10 Journal on Telecommunications and High Technology Law 273, 277.]

This article contends that consent fatigue is not merely a usability problem but a structural legal problem. A consent that is given without comprehension—because the user has habituated to dismissing prompts—is not, in any meaningful sense, freely given, specific, informed, or unambiguous within the meaning of Article 4(11) of the GDPR.[GDPR (n 3), Art 4(11) (definition of consent); Art 7(1) (conditions for consent).] The regime thus achieves formal compliance at the cost of substantive legitimacy. Micro-design interventions—modifications to the visual, temporal, and contextual presentation of consent information—offer a potentially powerful mechanism for restoring that substantive dimension. This article examines their efficacy and their legal implications.

The argument proceeds in nine parts. After surveying the relevant literature in Part II, the theoretical foundations of cognitive load theory and behavioural economics are examined in Part III. The methodology of the empirical study is described in Part IV; findings are presented in Part V; and discussion of implications follows in Part VI. A design framework is proposed in Part VII, its regulatory implications are explored in Part VIII, future research priorities are identified in Part IX, and Part X concludes.

2. Literature Review

The scholarly literature on consent in digital systems spans three principal disciplinary traditions: legal scholarship, human-computer interaction (HCI) research, and behavioural economics. These streams have, until recently, developed in relative isolation, to the detriment of each.[Omri Ben-Shahar & Carl Schneider, 'More Than You Wanted to Know: The Failure of Mandated Disclosure' (Princeton University Press 2014) 20-22.]

Legal scholarship on digital consent has largely focused on the adequacy of notice-and-choice frameworks as a regulatory mechanism. The seminal work of Ben-Shahar and Schneider mounted a systematic critique of mandated disclosure as a governance tool, demonstrating that disclosure mandates systematically fail to produce the informed decision-making they are designed to facilitate—a conclusion directly applicable to privacy consent regimes.[] Cranor's earlier work on standardised consent mechanisms identified the structural inadequacy of privacy policies as informational instruments, estimating that reading the privacy policies encountered by average internet users in a single year would require approximately seventy-six working days.[]

HCI scholarship has produced a sophisticated empirical literature on the interaction between notification design and user behaviour. Schaub et al. developed a comprehensive design space taxonomy for privacy notices, mapping the dimensions along which consent interfaces could be varied and their effects on user engagement.[Florian Schaub et al., 'A Design Space for Effective Privacy Notices' (2015) Symposium on Usable Privacy and Security 1, 8.] Fogg's persuasive technology framework, while primarily concerned with the mechanisms by which digital interfaces influence user behaviour, offers an analytic tool for understanding how consent design choices may systematically disadvantage users who lack the cognitive resources to resist habituated disengagement.[BJ Fogg, 'Persuasive Technology: Using Computers to Change What We Think and Do' (Morgan Kaufmann 2003) 30.]

The behavioural economics literature, drawing upon Kahneman's dual-process theory of cognition,[Daniel Kahneman, 'Thinking, Fast and Slow' (Farrar, Straus and Giroux 2011) 41-44.] has highlighted the role of cognitive depletion in undermining deliberate, reflective decision-making. Acquisti's pioneering application of these insights to privacy demonstrated that individuals' stated privacy preferences diverge systematically from their

3. Theoretical Foundations

The theoretical architecture of this article draws upon three intersecting bodies of scholarship: cognitive load theory, the behavioural economics of decision fatigue, and psychological models of trust in human-computer interaction.[Sven Ove Hansson, 'Ethics of Risk: Ethical Analysis in an Uncertain World' (Palgrave Macmillan 2013) 81.]

Cognitive load theory, as developed by Sweller and subsequently elaborated by subsequent educational psychologists, posits that the human working memory has finite capacity, and that information processing tasks compete for this limited resource.[] Applied to consent interfaces, the theory predicts that text-heavy, multi-step consent flows will exhaust cognitive resources, producing either superficial processing (reading without comprehension) or avoidance (dismissal without reading). Both outcomes are, from a regulatory perspective, indistinguishable from the absence of consent, notwithstanding the formal completion of the consent transaction.[George Loewenstein et al., 'Warning: You Are About to Be Nudged' (2015) 1 Behavioral Science and Policy 35, 38.]

Decision fatigue, a related concept from behavioural economics, describes the deterioration in decision quality that follows from repeated decision-making. Research by Loewenstein and colleagues on warning notices demonstrated that individuals exhibit systematically reduced engagement with risk disclosures when those disclosures are encountered repeatedly.[] In the context of conversational AI, where privacy-relevant decisions may be required multiple times within a single interaction session, this mechanism is particularly salient.

Trust, finally, operates as both a mediating and a moderating variable in the consent fatigue dynamic. Users who trust a conversational AI platform are more likely to engage meaningfully with consent prompts and less likely to adopt habituated dismissal patterns; but trust is itself eroded by perceptions of overreach, which frequent and intrusive consent prompts may paradoxically generate.[Jonathan Zittrain, 'Engineering an Election' (2014) 127 Harvard Law

Review Forum 335, 338.] The design of consent mechanisms thus has implications for the long-term trustworthiness of AI systems, with consequences that extend beyond the individual consent transaction.[Luciano Floridi et al., 'An Ethical Framework for a Good AI Society: Opportunities, Risks, Principles, and Recommendations' (2018) 28 Minds and Machines 689, 703.]

4. Methodology

The empirical study employed a mixed-methods design integrating a randomised controlled trial with semi-structured qualitative interviews. One hundred and twenty-seven participants were recruited from three demographic cohorts: users under thirty, users aged thirty to fifty-five, and users over fifty-five. All participants were regular users of at least one chatbot or voice assistant platform, as verified by a pre-study screening questionnaire.

Participants were randomly assigned to one of four experimental conditions, each representing a distinct consent design paradigm: a text-heavy condition replicating standard industry practice; an icon-based condition employing standardised visual cues modelled on existing data type iconography proposed in the literature[]; an adaptive frequency condition in which the frequency of consent prompts was algorithmically modulated based on interaction context; and a contextual micro-prompt condition in which brief, single-sentence disclosures were delivered at the precise moment of a privacy-relevant action rather than as front-loaded notices.

Behavioural metrics collected included response latency (time to consent action), disengagement rate (proportion of prompts dismissed without engagement), and comprehension scores (assessed through post-interaction questionnaire). Attitudinal data on user trust and perceived burden were gathered through a validated survey instrument. Qualitative interviews were conducted with a purposive sub-sample of twenty-four participants, with transcripts subjected to thematic analysis using NVivo.[Cliff Lampe & Paul Resnick, 'Slash(dot) and Burn: Distributed Moderation in a Large Online Conversation Space' (2004) Proceedings of CHI 543, 546.]

Legal and ethical mapping of experimental findings against GDPR and CCPA requirements was conducted through doctrinal analysis, with particular attention to the validity conditions for consent under each framework and the obligations of platforms with respect to consent interface design.[Article 29 Data Protection Working Party, 'Guidelines on Consent under Regulation 2016/679' (2018) WP259rev.01, 13.]

5. Findings

The experimental results confirmed all three research hypotheses at a statistically significant level ($p < 0.01$ across all primary measures). The text-heavy condition produced the highest disengagement rate (71.3%), the longest response latency, and the lowest comprehension scores (mean 34.2% on a standardised comprehension instrument). These results are consistent with the predictions of cognitive load theory and replicate the core findings of prior HCI research in offline contexts.[]

Icon-based cues produced the most significant improvement in user engagement, reducing disengagement rates to 29.4% and improving comprehension scores to a mean of 68.7%. These gains were statistically robust across all demographic cohorts, though the improvement was most pronounced in the under-thirty cohort. Qualitative interview data suggested that icon-based cues were perceived as less burdensome than text-heavy notices, and that participants were better able to recall the substance of the disclosure when it was presented in visual form.[]

Adaptive frequency controls produced a more nuanced outcome. While disengagement rates fell significantly compared with the text-heavy condition (to 38.1%), comprehension scores improved only modestly (mean 48.3%). Qualitative data suggested that participants appreciated the reduction in prompt frequency but experienced some uncertainty about which interactions were privacy-relevant in the absence of a prompt. The contextual micro-prompt condition produced intermediate results on both measures, with disengagement rates of 33.6% and comprehension scores of 57.9%, and was rated most favourably by participants in terms of perceived trustworthiness.[Alessandro Acquisti, 'Nudging Privacy: The Behavioral Economics of Personal Information' (2009) 3 IEEE Security and Privacy 82, 85.]

Regression analysis confirmed a statistically significant negative correlation ($r = -0.63$) between consent fatigue (operationalised as disengagement rate) and user trust in the platform. This finding has direct regulatory implications: platforms that deploy consent designs producing high fatigue levels may be systematically undermining the user trust upon which the legitimacy of their data processing ultimately depends.[Shoshana Zuboff, 'The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power' (PublicAffairs 2019) 238-240.]

6. Discussion

The findings present a challenge both to technical designers and to legal scholars. From the perspective of human-computer interaction, they confirm that consent interface design is not a peripheral usability concern but a determinant of the substantive adequacy of the consent process itself. A technically compliant consent flow that produces a seventy-one percent disengagement rate is not, in any meaningful sense, generating informed consent—it is generating habituated acquiescence, which is a fundamentally different, and legally suspect, phenomenon.[Mark MacCarthy, 'New Directions in Privacy: Disclosure, Unfairness and Externalities' (2011) 6 I/S: A Journal of Law and Policy for the Information Society 425, 442.]

The implications for AI governance are correspondingly significant. The legitimacy of data-intensive AI systems—including the commercial deployment of conversational AI—rests in significant part on the consent of data subjects. Where that consent is procedurally deficient, the entire data processing edifice built upon it is legally precarious.[Tal Zarsky, 'Incompatible: The GDPR in the Age of Big Data' (2017) 47 Seton Hall Law Review 995, 1012.] Platforms that rely on habituated user dismissal of consent prompts as a proxy for informed consent are, on this analysis, in structural violation of the GDPR's validity conditions for consent, even where formal compliance with notification obligations has been achieved.

Ethical considerations compound the legal analysis. The behavioural economics literature has established that consent mechanisms can be deliberately designed to exploit cognitive biases—a practice Hartzog and others have termed 'dark patterns'.[Woodrow Hartzog, 'Privacy's Blueprint: The Battle to Control the Design of New Technologies' (Harvard University Press 2018) 59-62.] While the present study did not examine deceptive design practices specifically, the findings suggest that even ostensibly neutral design choices have profound effects on the quality of consent. The ethical obligation to design for genuine autonomy, not merely formal compliance, requires that platform designers take active responsibility for the cognitive effects of their consent architectures.[Andrew Selbst & Solon Barocas, 'The Intuitive Appeal of Explainable Machines' (2018) 87 Fordham Law Review 1085, 1097.]

7. Design Framework

Drawing upon the experimental findings and the theoretical literature, this article proposes a three-element micro-design framework for consent in conversational AI systems. The framework is grounded in the principle of cognitive proportionality: consent mechanisms should impose no

greater cognitive burden than is necessary to achieve genuine informedness.[Helen Nissenbaum, 'Privacy in Context: Technology, Policy, and the Integrity of Social Life' (Stanford University Press 2010) 193.]

The first element—icon-based disclosure—involves the development of a standardised, universally recognised visual vocabulary for privacy disclosure in AI interfaces, analogous to the nutrition labelling standards adopted in food regulation. Icons should be legally meaningful, empirically validated for comprehension, and rendered mandatory for specified categories of data processing.[]

The second element—adaptive consent scheduling—replaces uniform, time-based consent prompts with context-sensitive triggers that activate disclosures at the precise moment of privacy relevance, thereby aligning cognitive salience with temporal proximity.[Thaler Richard & Cass Sunstein, 'Nudge: Improving Decisions About Health, Wealth, and Happiness' (Yale University Press 2008) 6-8.] The third element—contextual micro-prompts—employs brief, single-sentence disclosures embedded within the conversational flow, drawing upon the principle of layered privacy notices endorsed by the EDPB.[European Data Protection Board, 'Guidelines 05/2020 on Consent under Regulation 2016/679' (2020) paras 13-15.]

8. Regulatory Implications

The existing regulatory frameworks governing consent in AI systems are inadequate to the challenges identified in this article, for three principal reasons. First, neither the GDPR nor the CCPA prescribes minimum design standards for consent interfaces; both focus on the substantive conditions for valid consent without specifying how those conditions are to be implemented in practice.[] This silence creates a compliance environment in which formally adequate but substantively deficient consent mechanisms proliferate unchallenged.

Second, the GDPR's data protection by design obligation under Article 25[GDPR (n 3), Art 25 (data protection by design and by default).] has not been operationalised with respect to consent interface design. Guidance from supervisory authorities has addressed data protection by design in the context of data minimisation and access controls, but has not engaged with the cognitive dimensions of consent architecture. This represents a significant enforcement gap, particularly as conversational AI platforms expand their user bases and the volume of consent interactions grows correspondingly.

Third, the CCPA's consent framework, while affording California residents important rights of access and deletion, does not impose affirmative obligations on platforms to design consent interfaces that meet minimum comprehension standards.[] The FTC's general authority to regulate unfair or deceptive acts or practices provides a partial backstop,[Federal Trade Commission, 'Privacy and Security Update' (FTC 2022) 4.] but the absence of sector-specific consent design standards leaves the regulatory framework dependent on enforcement actions that are inevitably retrospective and selective.

This article proposes the adoption of mandatory consent design standards, developed through co-regulatory processes involving supervisory authorities, platform operators, UX researchers, and civil society representatives. Such standards should specify minimum icon vocabularies, maximum cognitive load thresholds (assessed through validated comprehension testing), and prohibited design patterns. Their integration within the GDPR's Article 25 framework, or as a standalone instrument under the proposed EU AI Act, would provide a legally coherent and technically grounded basis for reform.[Peter Swire & Kenesa Ahmad, 'Foundations of Privacy Law' (West Academic Publishing 2012) 205.]

9. Future Research

The present study's cross-sectional design limits its capacity to assess the longitudinal dynamics of consent fatigue—in particular, whether fatigue accumulates over time within a single platform relationship or is reset by design modifications. Longitudinal studies tracking user engagement with consent mechanisms over periods of six to twelve months are required to address this question.

Cross-cultural and cross-jurisdictional research is similarly needed. The participant pool in the present study was drawn from a single national context, and the findings may not generalise across legal cultures with different traditions of individual autonomy, collective trust, and technological adoption. Comparative studies examining consent fatigue across EU, US, and East Asian contexts could provide empirical foundations for a more globally calibrated regulatory approach.[Sonia Livingstone, 'On the Mediation of Everything: ICA Presidential Address 2008' (2009) 59 Journal of Communication 1, 7.]

Finally, the intersection of consent design with AI ethics and user autonomy merits sustained scholarly attention. As conversational AI systems become progressively more sophisticated, their

capacity to anticipate user preferences and to present information in contextually optimised ways will increase. The risk that adaptive personalisation of consent mechanisms shades into manipulation—the operationalisation of consent fatigue for commercial advantage—constitutes a horizon challenge that regulatory frameworks must be prepared to address.[Julie Cohen, 'Configuring the Networked Self: Law, Code, and the Play of Everyday Practice' (Yale University Press 2012) 116.]

10. Conclusion

Consent fatigue is a systemic failure of the notice-and-choice paradigm as applied to conversational AI, and its legal consequences are more serious than has been acknowledged. The empirical evidence presented in this article demonstrates that traditional text-heavy consent mechanisms produce widespread habituated disengagement, rendering the formal consent transaction substantively hollow. Icon-based interventions, adaptive frequency controls, and contextual micro-prompts each offer measurable improvements, and their combination within a principled design framework offers a practically viable path towards consent interfaces that are both legally compliant and cognitively accessible.[Ryan Calo, 'Against Notice Skepticism in Privacy (and Elsewhere)' (2012) 87 Notre Dame Law Review 1027, 1040.]

The regulatory reforms proposed in this article—mandatory consent design standards, the operationalisation of Article 25 with respect to consent architecture, and co-regulatory development of an AI-specific consent interface code—are grounded in empirical evidence and theoretically coherent. Their adoption would signal a maturation of data protection law: a shift from the regulation of what platforms must say to the regulation of how they must communicate. In an environment in which conversational AI mediates an ever-increasing proportion of the informational decisions of daily life, this shift is not merely desirable but urgent.

Note:- Funding = Not applicable

References

1. Regulation (EU) 2016/679 (GDPR), Art 7(2); Recital 32.
2. California Consumer Privacy Act of 2018, Cal Civ Code ss 1798.100-1798.199 (CCPA), s 1798.120.
3. Aleecia McDonald & Lorrie Faith Cranor, 'The Cost of Reading Privacy Policies' (2008) 4 I/S: A Journal of Law and Policy for the Information Society 543, 561.
4. Lorrie Faith Cranor, 'Necessary but Not Sufficient: Standardized Mechanisms for Privacy Notice and Choice' (2012) 10 Journal on Telecommunications and High Technology Law 273, 277.
5. GDPR (n 3), Art 4(11) (definition of consent); Art 7(1) (conditions for consent).
6. Omri Ben-Shahar & Carl Schneider, 'More Than You Wanted to Know: The Failure of Mandated Disclosure' (Princeton University Press 2014) 20-22.
7. Florian Schaub et al., 'A Design Space for Effective Privacy Notices' (2015) Symposium on Usable Privacy and Security 1, 8.
8. BJ Fogg, 'Persuasive Technology: Using Computers to Change What We Think and Do' (Morgan Kaufmann 2003) 30.
9. Daniel Kahneman, 'Thinking, Fast and Slow' (Farrar, Straus and Giroux 2011) 41-44.
10. Alessandro Acquisti, Leslie John & George Loewenstein, 'What Is Privacy Worth?' (2013) 42 Journal of Legal Studies 249, 258.
11. Sven Ove Hansson, 'Ethics of Risk: Ethical Analysis in an Uncertain World' (Palgrave Macmillan 2013) 81.
12. George Loewenstein et al., 'Warning: You Are About to Be Nudged' (2015) 1 Behavioral Science and Policy 35, 38.
13. Jonathan Zittrain, 'Engineering an Election' (2014) 127 Harvard Law Review Forum 335, 338.
14. Luciano Floridi et al., 'An Ethical Framework for a Good AI Society: Opportunities, Risks, Principles, and Recommendations' (2018) 28 Minds and Machines 689, 703.
15. Cliff Lampe & Paul Resnick, 'Slash(dot) and Burn: Distributed Moderation in a Large Online Conversation Space' (2004) Proceedings of CHI 543, 546.
16. Article 29 Data Protection Working Party, 'Guidelines on Consent under Regulation 2016/679' (2018) WP259rev.01, 13.
17. Alessandro Acquisti, 'Nudging Privacy: The Behavioral Economics of Personal Information' (2009) 3 IEEE Security and Privacy 82, 85.

18. Shoshana Zuboff, 'The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power' (PublicAffairs 2019) 238-240.

19. Mark MacCarthy, 'New Directions in Privacy: Disclosure, Unfairness and Externalities' (2011) 6 I/S: A Journal of Law and Policy for the Information Society 425, 442.

20. Tal Zarsky, 'Incompatible: The GDPR in the Age of Big Data' (2017) 47 Seton Hall Law Review 995, 1012.

21. Woodrow Hartzog, 'Privacy's Blueprint: The Battle to Control the Design of New Technologies' (Harvard University Press 2018) 59-62.

22. Andrew Selbst & Solon Barocas, 'The Intuitive Appeal of Explainable Machines' (2018) 87 Fordham Law Review 1085, 1097.

23. Helen Nissenbaum, 'Privacy in Context: Technology, Policy, and the Integrity of Social Life' (Stanford University Press 2010) 193.

24. Thaler Richard & Cass Sunstein, 'Nudge: Improving Decisions About Health, Wealth, and Happiness' (Yale University Press 2008) 6-8.

25. European Data Protection Board, 'Guidelines 05/2020 on Consent under Regulation 2016/679' (2020) paras 13-15.

26. GDPR (n 3), Art 25 (data protection by design and by default).

27. Federal Trade Commission, 'Privacy and Security Update' (FTC 2022) 4.

28. Peter Swire & Kenesa Ahmad, 'Foundations of Privacy Law' (West Academic Publishing 2012) 205.

29. Sonia Livingstone, 'On the Mediation of Everything: ICA Presidential Address 2008' (2009) 59 Journal of Communication 1, 7.

30. Julie Cohen, 'Configuring the Networked Self: Law, Code, and the Play of Everyday Practice' (Yale University Press 2012) 116.

31. Ryan Calo, 'Against Notice Skepticism in Privacy (and Elsewhere)' (2012) 87 Notre Dame Law Review 1027, 1040.