




Written Evidence on Transparency and Information Integrity under Article 50 of the EU AI Act

June 2026

Professor Ann Kristin Glenster and Peter Bannister



This report presents the evidence submitted by Professor Ann Kristin Glenster and Peter Bannister, LOGOS at the University of Cambridge, to the European Commission's Stakeholder consultation on the draft Guidelines on the transparency requirements for certain AI systems under Article 50 of the AI Act on 3 June 2026.

DOI: TBC

www.logosc.cambridge.org

© LOGOS 2026

Executive Summary

LOGOS argues that Article 50 transparency should function as a safeguard for human recognition, judgement and information integrity, not as a narrow notice obligation. In this evidence, transparency has three connected functions. It is recognitional because it helps natural persons identify AI interaction and AI-generated or manipulated content. It is epistemic because it helps them assess the source, status and communicative role of information. It is systemic because it helps preserve reliable signals across AI-mediated information chains. The effectiveness of Article 50 will depend on whether these functions are supported in practice at the point of interaction or exposure.¹

The central risk is that Article 50 transparency is implemented formally but not functionally. A disclosure may exist but be ignored, misunderstood, hidden in metadata, stripped from content, lost when content is reposted or transformed, or too limited to help a person understand the nature of the AI-mediated interaction. In such cases, transparency may satisfy a formal requirement while failing to support human recognition, calibrated trust or reliable assessment of information.²

These six recommendations follow from the recognitional, epistemic and systemic functions that Article 50 transparency must serve. First, the final Guidelines should clarify the threshold concepts that determine when Article 50 obligations apply, including “obvious”, “substantially alter”, “editorial control” and “public interest”. Secondly, they should distinguish individual-facing disclosure from wider information-chain integrity. Thirdly, they should clarify the role of GPAI model-level capabilities in downstream Article 50 compliance. Fourthly, they should strengthen guidance on the substance, design and evidencing of disclosures for interactive AI systems. Fifthly, they should clarify the distinction between standard editing, non-substantial alteration and meaningful AI manipulation. Finally, they should ensure that technical transparency signals are meaningfully available to natural persons at the point of exposure.

The distinctive contribution of this evidence is to treat Article 50 transparency as a practical condition for accountable AI-mediated communication. Transparency is not only a matter of telling people that AI is present. It is also a matter of preserving reliable signals about source, status, alteration, mediation and provenance so that natural persons can understand the communicative environment in which they are acting.

Recommendations

Recommendation 1. Clarify the threshold concepts that determine when Article 50 obligations apply.

The final Guidelines should make clearer that concepts such as “obvious”, “substantially alter”, “editorial control” and “public interest” operate as threshold tests rather than merely descriptive terms. Providers and deployers need practical guidance on how these thresholds should be assessed, documented and evidenced in context.

Recommendation 2. Distinguish individual-facing disclosure from wider information-chain integrity.

The Guidelines should not assume that informing individuals about AI-origin will, by itself, enable meaningful user choice or protect trust in the information ecosystem. Information asymmetries, cognitive overload and limited user choice may reduce the practical effect of disclosure. The final Guidelines should

¹ Glenster, A.K. (2025) *Surveillance Capitalism, Deceptive Designs, and Epistemic Rights*. Technology & Human Rights Discussion Paper. Cambridge, MA: Carr-Ryan Center for Human Rights, Harvard Kennedy School.

² Ananny, M. & Crawford, K. (2018) 'Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability', *New Media & Society*, 20(3), 973–989.

therefore distinguish user-facing transparency from wider measures that preserve reliable signals across the information chain.

Recommendation 3. Clarify the role of GPAI model-level capabilities in downstream Article 50 compliance.

The final Guidelines should avoid creating the impression that model-level transparency measures are outside the practical compliance chain. Even where Article 50 obligations apply to downstream AI systems, effective compliance may depend on upstream capabilities that support marking, detection, provenance, documentation or disclosure. Further guidance is needed on how responsibility should be allocated where downstream providers or deployers rely on such capabilities.

Recommendation 4. Strengthen guidance on the substance, design and evidencing of disclosures for interactive AI systems.

The final Guidelines should clarify whether informing natural persons of the “artificial, non-human nature” of an interaction requires only a statement that the interaction is with an AI system, or whether further information is required about the system’s role, function, limits, provider, deployment context, use of data or possible profiling. Disclosure design should be assessed by reference to whether the information is reasonably noticeable, understandable and meaningful to natural persons in the relevant context of use.

Recommendation 5. Clarify the distinction between standard editing, non-substantial alteration and meaningful AI manipulation.

The final Guidelines should provide further practical guidance on when AI-assisted changes to style, fluency, tone, structure, format, presentation or emphasis amount to a substantial alteration of content or semantics. This is especially important because many commonly used AI systems alter how information is presented without clearly changing its factual content, while still potentially affecting meaning, impression or interpretation.

Recommendation 6. Ensure that technical transparency signals are meaningfully available to natural persons at the point of exposure.

Machine-readable marking, metadata and provenance signals may support detection and information-chain integrity, but they are not equivalent to human-facing transparency. Where Article 50 requires information to be provided to natural persons, the Guidelines should clarify how technical signals should be surfaced in a clear, distinguishable and context-sensitive manner, including where content is reposted, embedded, summarised, transformed or encountered in high-volume information environments.

About this Evidence

1. This written evidence develops the response submitted by LOGOS to the European Commission’s consultation on the draft Guidelines on the transparency requirements for certain AI systems under Article 50 of the EU AI Act on 3 June 2026. It is intended as a fuller published version of that response, setting out the underlying argument, recommendations and evidential basis for LOGOS’s position.
2. Article 50 establishes transparency obligations for specified AI systems and AI-generated or manipulated content, including obligations relating to interaction with AI systems, marking and detection of AI-generated or manipulated content, emotion recognition and biometric categorisation systems, and labelling of deep fakes and certain AI-generated or manipulated text published to inform the public on matters of public interest.³

³ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) [2024] OJ L, 2024/1689, Article 50.

3. LOGOS is a research-to-impact initiative at the University of Cambridge with a mission to conduct and use rigorous research to protect and promote epistemic rights; information integrity and data supply chains; and reliable, independent and secure global digital infrastructure. Its aim is to help ensure that emerging technologies benefit the advancement of human knowledge and cognitive agency.

4. This evidence focuses on the general interpretation of Article 50 transparency obligations, transparency for interactive AI systems, marking and detection of AI-generated or manipulated content, labelling of deep fakes and certain public-interest text, and horizontal requirements for clear, meaningful and accessible information. Article 50(3), on emotion recognition and biometric categorisation systems, is not treated as a standalone section because LOGOS's detailed consultation comments focused primarily on the other Article 50 obligations. The implications for Article 50(3) are addressed through the report's wider discussion of meaningful, timely and human-facing transparency.

5. LOGOS welcomes the Commission's effort to provide practical guidance on the implementation of Article 50. Clear Guidelines are necessary to support consistent interpretation, practical compliance and effective enforcement across the Union. However, the draft Guidelines should be strengthened where they use concepts that determine whether transparency obligations apply, how they should be implemented, or whether exceptions may be relied upon.⁴

6. The central argument of this evidence is that Article 50 transparency should not be treated only as a formal notice obligation. In AI-mediated environments, transparency should operate as a safeguard for human recognition, judgement and trust. It should enable natural persons to recognise when they are interacting with AI systems, when content has been generated or manipulated by AI, and when AI-mediated communication may affect the conditions under which information is assessed, relied upon and circulated.

Article 50 Transparency and Information Integrity

7. Transparency under Article 50 is not a single mechanism. The obligations in Article 50 include different forms of transparency, including disclosure to natural persons, marking of AI-generated or manipulated content, technical measures for detection, labelling of deep fakes and certain public-interest text, and horizontal requirements that information be provided in a clear and distinguishable manner. These mechanisms may overlap, but they do not perform the same function.⁵

8. Article 50 transparency has a recognitional function. Natural persons should be able to recognise when they are interacting with an AI system or encountering AI-generated or manipulated content. This function is central to reducing risks of impersonation, deception and anthropomorphisation, and to supporting trust in the information ecosystem.⁶

9. Article 50 transparency also has an epistemic function. It helps natural persons assess the source, status and communicative role of information. In AI-mediated environments, a person may need to know whether AI generated content, substantially altered it, summarised it, mediated an interaction, acted on behalf of another person, or supplied a technical signal that is not visible at the point of

⁴ European Commission, Directorate-General for Communications Networks, Content and Technology, Artificial Intelligence Office, *Draft Guidelines on the implementation of the transparency obligations for certain AI systems under Article 50 of Regulation (EU) 2024/1689 (the 'AI Act')*, draft for stakeholder consultation, May 2026.

⁵ See note 4 above (draft Guidelines, §2.1)

⁶ See note 10 below on the NIST Human-AI Configuration risk category.

exposure. Transparency is therefore connected to the practical conditions under which natural persons can assess information and decide whether, and how, to rely on it.

10. Article 50 transparency also has a systemic function. It should help preserve reliable signals across the information chain. A disclosure shown to one user at one point in a system may not remain visible, meaningful or attached when content is copied, reposted, embedded, summarised, transformed or circulated through high-volume information environments. For that reason, Article 50 transparency should be understood not only as individual-facing disclosure, but also as part of the wider infrastructure of information integrity.⁷

11. These three functions explain why Article 50 transparency should be assessed functionally rather than formally. A disclosure, label or technical signal is effective only if it supports recognition, informed assessment and the preservation of reliable signals in the context in which the natural person encounters the system or content.

Threshold Concepts Requiring Clarification

12. Several concepts in the draft Guidelines operate as legal and practical gateways to obligation. This is particularly important in relation to “obvious” under Article 50(1), “substantially alter” under Article 50(2), “editorial control” under Article 50(4), and “public interest” under Article 50(4). These terms determine whether obligations apply, whether exceptions may be relied upon, and how providers and deployers should assess their responsibilities.

13. The importance of these concepts lies not only in their interpretation, but also in their practical consequences. If a threshold is interpreted too broadly, transparency obligations may be avoided in circumstances where natural persons still need information to understand an AI-mediated interaction or communication. If interpreted too narrowly, obligations may be triggered in ways that are difficult to implement or that add little practical value for natural persons. The final Guidelines should therefore treat these concepts as operational thresholds requiring context-specific assessment, documentation and evidencing.

14. Cross-cutting guidance is needed on how threshold assessments should be made. Such assessments should not rest only on provider intention, formal availability of information, or general assumptions about what users ought to understand. They should take account of the relevant context of use, the likely perspective of natural persons, the system’s design and modality, the communicative effect of the content or interaction, and the practical consequences of relying on an exception.

15. This cross-cutting guidance would support the more specific analysis required under each Article 50 obligation. The obviousness exception under Article 50(1), the distinction between standard editing and substantial alteration under Article 50(2), and the concepts of editorial control and public interest under Article 50(4) raise different legal and practical questions. However, they share a common implementation problem: each requires providers or deployers to justify why a transparency obligation does, or does not, apply in context.

16. Providers and deployers should be able to evidence these threshold assessments. Where a provider or deployer concludes that an interaction is obvious, that content has not been substantially altered, that an actor lacks editorial control, or that content is outside the scope of public interest,

⁷ Wardle, C. & Derakhshan, H. (2017) *Information Disorder: Toward an Interdisciplinary Framework for Research and Policymaking*. Council of Europe report DGI(2017)09. Strasbourg: Council of Europe.

the final Guidelines should clarify how that assessment should be documented. This would support consistent compliance, supervisory oversight and meaningful accountability.

Information-Chain Integrity and Point-of-Exposure Transparency

17. The distinction between individual-facing disclosure and information-chain integrity is central to this evidence. A disclosure to one person at one moment is not the same as a reliable transparency signal that remains available as content moves through systems, platforms, interfaces and publication workflows. Article 50 will be more effective if the Guidelines recognise both dimensions.

18. In practice, a transparency signal may be lost before the natural person most affected by it encounters the content. A label may be removed when a file is compressed, a metadata signal may not be surfaced by a platform, or a visible disclosure may be detached when content is screenshotted, embedded or summarised. The result is that transparency exists at the point of origin but fails at the point of exposure.⁸

19. The final Guidelines should therefore address preservation as well as provision. Providers and deployers need guidance not only on when to mark, label or disclose, but also on how to preserve and surface transparency information when content is transmitted, transformed or republished. This is especially important where AI-generated or manipulated content affects personal identity, public-interest communication, or the perceived authenticity of audio, visual or textual information.

20. Point-of-exposure transparency should be treated as a practical test of effectiveness. Where Article 50 requires information to be provided to natural persons, the relevant question is whether the person can reasonably perceive and understand the transparency information in the context in which they encounter the content or system. Technical marking, metadata and provenance signals can support this aim, but they do not replace human-facing transparency.

Downstream Compliance, Dissemination and GPAI Model-Level Capabilities

21. Article 50 obligations operate across a practical compliance chain. Even where an Article 50 obligation applies to a downstream AI system, effective compliance may depend on capabilities, documentation or technical measures that originate upstream. This is particularly relevant to marking, detection, provenance and disclosure, where downstream providers or deployers may not be able to implement meaningful transparency without support from upstream model-level capabilities.

22. The statement that Article 50 does not explicitly apply to general-purpose AI models should therefore be framed carefully. It should not create the impression that GPAI model-level transparency measures are irrelevant to the practical implementation of Article 50. Where downstream Article 50 obligations depend in practice on model-level capabilities, the Guidelines should explain how providers and deployers should approach that dependency.

23. Further guidance is needed on the allocation of responsibility between actors in the AI value chain. Providers and deployers need to understand what they are expected to do where transparency depends on information, technical affordances or documentation supplied by another actor. The Guidelines should therefore clarify how responsibility should be allocated where downstream

⁸ Partnership on AI (2023) *PAI's Responsible Practices for Synthetic Media: A Framework for Collective Action*. San Francisco: Partnership on AI.

providers or deployers rely on upstream capabilities for marking, detection, provenance or disclosure.

24. The Guidelines should also clarify the practical status of actors whose role is limited to disseminating or transmitting third-party AI-generated or manipulated content. If such actors are not deployers but are encouraged to preserve marking and labelling or to take appropriate measures, the Guidelines should explain the practical meaning of that encouragement. In particular, it should clarify whether this is a best-practice expectation, a recommended risk-mitigation measure, or a factor relevant to the preservation of reliable transparency signals across the information chain.⁹

Article 50(1): Transparency for Interactive AI Systems

25. Article 50(1) requires particular care because interactive AI systems may affect how natural persons understand the nature of an interaction. The distinction between direct interaction with an AI system and human interaction that is mediated, assisted or partly shaped by an AI system should therefore be clarified. Natural persons may encounter interactions that are not solely machine interactions, but are nevertheless significantly structured by AI.

26. Further guidance is especially needed where AI systems act on behalf of natural persons. An AI agent may communicate, filter, negotiate, recommend, draft or respond in a way that represents or mediates a human actor. For example, an automated assistant may appear to speak on behalf of a person, institution or service while materially shaping the content and timing of the exchange. The Guidelines should clarify when such interactions fall within Article 50(1), and what information must be provided to the natural person interacting with the system.

27. The concept of avoiding over-reliance also requires clarification. Over-reliance is not resolved simply by informing a person that a system is AI. Natural persons may still place excessive trust in outputs or interactions where a system appears fluent, authoritative, personalised, human-like or contextually competent. Disclosure design should therefore be assessed by reference to whether it helps natural persons understand the nature and limits of the system in the specific context of use.¹⁰

28. The minimum substance of disclosure under Article 50(1) should be clarified. The Guidelines refer to informing natural persons of the artificial, non-human nature of the interacting counterpart. However, it is not always clear whether this requires only a statement that the interaction is with an AI system, or whether further information is required about the kind of AI system, its role, function, limits, provider, deployment context, use of data or possible profiling of personal data. The final Guidelines should clarify when such additional information is required for the disclosure to be meaningful.

29. This is particularly important where interactive AI systems use conversational, human-like or agentic design. A voice assistant, avatar or emotionally responsive chatbot may encourage users to interpret the interaction through social or relational expectations. In such contexts, disclosure should not only identify the system as AI but should help natural persons understand the nature of the interaction and avoid misinterpreting the system's capacities, authority or relationship to the user.¹¹

⁹ Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market for Digital Services (Digital Services Act), OJ L 277, 27.10.2022, p. 1.

¹⁰ National Institute of Standards and Technology (2024) *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*. NIST AI 600-1. Gaithersburg, MD: U.S. Department of Commerce.

¹¹ See note 10, particularly on emotional entanglement.

30. Disclosure format should also be assessed in context. Information should be clear, accessible and adapted to the modality of interaction. A text notice may be insufficient where the interaction is primarily voice-based, visual, immersive, embodied or multimodal. Providers and deployers should therefore determine the appropriate form, timing and modality of disclosure, including how disclosures should be adapted for users with accessibility needs.

31. The relationship between disclosure effectiveness, disruption, friction and habituation requires further guidance. The concern that disclosures may cause friction, or that users may become habituated to notices, should not be interpreted in a way that weakens Article 50(1). Rather, it should encourage better disclosure design. Transparency information should remain noticeable, understandable and proportionate to the risks and context of the interaction.

32. The obviousness exception should be accompanied by clearer evidencing expectations. Where a provider or deployer relies on the claim that interaction with an AI system is obvious from the circumstances and context of use, the Guidelines should clarify how that assessment should be made and documented. Evidence should relate to the perspective of natural persons in the relevant context, rather than to provider assumptions about what users ought to understand.

Article 50(2): Marking and Detection of AI-Generated or Manipulated Content

33. Article 50(2) raises distinct questions because marking and detection obligations concern not only what natural persons are told, but also how AI-generated or manipulated content remains identifiable across systems and contexts. The relationship between marking, detection, metadata, provenance signals and human-facing information should be clarified. Technical signals may support transparency, but they do not necessarily amount to meaningful information for natural persons where Article 50 requires human-facing disclosure.

34. The scope of content “intended to be directly perceived by natural persons” would benefit from further clarification. There may be cases where AI-generated or AI-manipulated content is not directly perceivable by natural persons but is still intended to influence an information process. The Guidelines should explain whether such content falls outside Article 50(2) because it is not intended for direct human perception, even where it may affect decisions, rankings, assessments or other forms of information processing.

35. Hidden or machine-facing content illustrates the difficulty. For example, hidden text inserted into CVs, job applications or other documents may be intended for automated processing rather than human perception. Such content may not be directly perceived by natural persons, but it may still affect the processing, ranking or interpretation of information. The Guidelines should clarify whether and how Article 50(2) applies to such cases, or whether other provisions or governance measures are expected to address them.

36. The exception for standard editing and non-substantial alteration remains difficult to apply. The Guidelines indicate that standard editing may include corrections to technical, grammatical or linguistic errors, layout, presentation, formatting, accessibility requirements or sectoral practices. They also indicate that editing goes beyond standard editing where content is changed in a material way affecting meaning, style or intent. Further practical guidance is needed on how this threshold should be applied in ordinary uses of AI systems.

37. The difficulty is especially acute because many AI tools alter style, fluency, tone, structure, presentation or format without clearly changing semantic content. A text may retain broadly the same factual claims while acquiring a different emphasis, register or persuasive effect. In some contexts, these changes may affect how the information is understood or trusted. The final

Guidelines should therefore provide examples that distinguish minor standard editing from changes that materially alter meaning, style, intent or communicative effect.

38. Providers should be able to conduct and document case-specific assessments of substantial alteration. A general statement that a system is used for editing, summarisation, enhancement or formatting may not be sufficient. Providers should be able to explain how they assess whether a system substantially alters input data or semantics, and how that assessment takes account of the type of content, modality, deployment context and likely effect on natural persons.

39. Implementation cost should not be read as negating the obligations under Article 50(2). Cost may be relevant to the selection of appropriate technical measures, proportionality and implementation choices, but it should not determine whether the underlying transparency obligation exists. Technical difficulty or implementation cost should not by itself justify treating marking or detection obligations as optional.

40. These preservation issues are addressed above as part of information-chain integrity. In the Article 50(2) context, initial marking and detection will have limited practical effect if transparency signals are not preserved, transmitted or surfaced when AI-generated or manipulated content moves across systems.

Article 50(4): Labelling of Deep Fakes and Certain Public-Interest Text

41. Article 50(4) requires careful interpretation because AI-generated or manipulated content may affect how natural persons assess authenticity, veracity and meaning. The Guidelines should clarify that deep fake content may include partial generation or manipulation of image, audio or video content. Where most of the content is not generated or manipulated, but a significant element is, disclosure may still be necessary if the other criteria under Article 50(4) are met.

42. Not every technical alteration has the same significance. Minor technical changes should be distinguished from manipulations that affect a person's assessment of authenticity, truthfulness or identity. Practical examples would help deployers distinguish ordinary technical changes from partial manipulations that materially affect how content is perceived or trusted.¹²

43. Public-interest communication is especially sensitive because it shapes the conditions under which natural persons form judgements about civic life, democratic participation, institutional trust and public debate. AI-generated or manipulated content in this context may not only misstate facts, but also alter emphasis, context, salience or perceived authority. Article 50(4) should therefore be applied in a way that protects the integrity of public-interest communication without unduly restricting lawful expression.¹³

44. The concept of text "published with the purpose of informing the public on matters of public interest" should be clarified. The Guidelines should avoid limiting public interest too narrowly to factual or verifiable information alone. Depending on context, public-interest communication may include cultural, artistic, satirical or analogous forms of expression that contribute to public understanding, democratic discourse or social critique. Practical examples would assist deployers in assessing when AI-generated or manipulated text falls within Article 50(4).

¹² Paris, B. & Donovan, J. (2019) *Deepfakes and Cheap Fakes: The Manipulation of Audio and Visual Evidence*. New York: Data & Society Research Institute.

¹³ Risse, M. (2023) *Political Theory of the Digital Age: Where Artificial Intelligence Might Take Us*. Cambridge: Cambridge University Press.

45. The disclosure obligation for AI-generated or manipulated text should also be clarified in relation to content, form and meaning. A text may contain accurate factual information while still changing the meaning or impression of the source through tone, emphasis, omissions, framing or structure. This is particularly relevant to AI-generated summaries of news, public-interest material or official information. The Guidelines should clarify whether generation or manipulation under Article 50(4) concerns the content of the text, the form in which it is presented, or both, and what the label should communicate in each case.

Article 50(5): Horizontal Requirements

46. Article 50(5) is central to the effectiveness of the transparency obligations because it concerns how information is provided to natural persons. The Guidelines should clarify whether information is sufficient merely because it has been made available, or whether providers and deployers should be able to evidence that it is reasonably noticeable, understandable and meaningful in the relevant context of use.

47. The distinction between machine-readable marking and human-facing disclosure is especially important. Machine-readable information, metadata or other technical signals may support detection, provenance and preservation of reliable signals across the information chain. However, where Article 50 requires information to be provided to natural persons, the Guidelines should clarify how such technical signals should be surfaced in a way that is clear and distinguishable to the person concerned.

48. This issue is particularly acute where AI-generated or manipulated content is encountered after being reposted, embedded, summarised, transformed or circulated through high-volume information environments. A disclosure may formally exist at the point of origin but no longer be visible, meaningful or distinguishable at the point of exposure. The Guidelines should therefore clarify how “first interaction or exposure” operates where content moves across systems, contexts and interfaces.

49. The accessibility of transparency information should also be assessed in context. Information should be provided in a form that natural persons can reasonably perceive and understand, taking account of context, modality, language, disability, age, device, interface and the likely conditions of exposure. This does not require a single disclosure format for all cases. It requires a context-sensitive assessment of whether the information is practically available to the natural person concerned.

50. Article 50(5) should therefore be understood as more than a formal presentation requirement. It is the provision that connects Article 50’s different transparency obligations to their practical effect. Without clear, distinguishable and context-sensitive presentation, disclosures, labels, markings and metadata may exist in formal or technical terms while failing to provide meaningful transparency to natural persons.

Conclusion

51. Article 50 transparency will only succeed if it is implemented as a meaningful safeguard for human recognition, judgement and information integrity. The final Guidelines should therefore clarify not only what information must be provided, but how transparency obligations should operate in the practical environments in which AI systems and AI-generated or manipulated content are encountered.

52. The next implementation challenge is practical. Providers, deployers, intermediaries and competent authorities will need workable evidencing standards, preservation practices and human-facing presentation rules. These operational conditions will determine whether Article 50 transparency functions as meaningful transparency or as a formal compliance layer.

53. LOGOS will continue to examine how AI-mediated information environments affect epistemic rights, information integrity and human cognitive agency. Article 50 can contribute to trustworthy AI governance if the final Guidelines treat transparency not as a formal compliance step, but as a practical condition for human recognition, reliable information and accountable AI-mediated communication.



LOGOS
CAMBRIDGE

www.logoscambidge.org