

PROTECTING THE THINKING PROCESS IN HUMAN-AI INTERACTION: THE *FORUM INTERSTITIALE* BETWEEN THOUGHT AND EXPRESSION

Caroline De Cock[*]

International human rights law protects freedom of thought absolutely and freedom of expression subject to limitations. The relevant established doctrine calls the two sides of that line the forum internum and the forum externum. However, the line was drawn when thinking left no trace, and with current technological developments, it no longer holds.

---

[*] Caroline De Cock has spent nearly thirty years working inside European Union digital policymaking across telecommunications, copyright, platform regulation, and artificial intelligence governance. She is the author of AI TOOLS, NOT GODS: WHY ARTIFICIAL INTELLIGENCE HYPE THREATENS GLOBAL GOVERNANCE—AND HOW TO FIX IT (2026) and writes on technology policy for Tech Policy Press. This Article forms part of a personal research project on freedom of thought in the age of artificial intelligence, with an accompanying website at https://foruminternum.org. The line of research began after she interviewed Jacob Mchangama, Executive Director of The Future of Free Speech, for the *AI lab TL;DR* podcast produced by information labs; his work on how models handle lawful but contested prompts set the direction of this Article's inquiry and is discussed in Part IV.C. The argument herein took further shape at the Summer School in Digital Human Rights held at Lund University in June 2026, organized by the Faculty of Law there and the Raoul Wallenberg Institute of Human Rights and Humanitarian Law. She is grateful to Jessica Almqvist and Marko Milanovic for convening and driving it, to Philippa Webb, Chantal Joris, Alexander Seger, Alberto Rinaldi, and Sue Anne Teo for sessions that shaped the questions this Article asks, and to the participants, whose experience across very different jurisdictions kept the discussion concrete. She thanks also Jeff Ubois, whose suggestion of the word interstitial gave the category its name. The views expressed here, and any errors, are the author's alone.

AI acknowledgment: This paper used AI tools to assist at various stages of research and organization. NotebookLM served to track source materials and manage bibliographic references to facilitate accurate citation. Claude Opus 5 was employed in the ideation phase for initial brainstorming and exploration, and served as a feedback provider during the revision of earlier drafts. Finally, Grammarly was used to improve grammar and style. All AI-assisted output underwent a process of critical integration, where it was rigorously evaluated and selectively incorporated by the author.

Within the span of six months in 2026, a United States court ordered that 20 million ChatGPT conversations be provided as evidence under ordinary discovery rules, two federal courts reached opposite conclusions on whether a litigant's exchanges with a chatbot were protected at all, and the Supreme Court held that location records kept by a third party remain constitutionally protected. Records of where a person's body had been were protected, while records of how their mind worked were not.

The law currently has no category for the search query, the exploratory prompt, the unfinished draft, or the confession made to a machine. The U.S. answer has been a quarrel about privilege: Courts refuse to grant AI exchanges between users and chatbots protection because the machine is not a professional, like a doctor or a lawyer, and the industry asks for a new privilege because it claims the model behaves like one. Both the courts and the industry assume that protection depends on what the person is using, i.e., what the machine is. This Article argues that it depends on what the person who uses the machine is doing.

It proposes the *forum interstitiale*, a third zone between the two classical fora, which can protect the thinking process under a standard stronger than ordinary expression law but short of absolute. Every candidate analogy to interactions with AI chatbots in existing law is tested herein, from search queries and confidential dialogue to protected reading, aggregated records, the European law on correspondence and data protection, private diaries and drafts, and neural data. Each fails for the same reason, namely that the protection it offers relates to a relationship, a disclosure, a recording, an addressee, or a data type, and never to the function of the act as thinking.

The interstitial zone this Article proposes only flows upwards. It lifts acts abandoned by the current binary approach into protection and displaces nothing from the absolute core of protected thought. The Article also sets out what acts belong in the liminal zone, five elements of the protection standard, the duties those elements impose on AI providers, and a map of individual AI interactions across the three fora. This classification determines what records of thought can be demanded, kept, filtered, and used against the person who was engaged in the process of thinking.

Introduction: When Three U.S. Court AI Decisions Create Uncertainty 224

I. The Classical Binary and Its Critics: Two Fora and a Gap in the Middle ... 229

A. The Architecture of the Two Fora 229

B. Three Lines of Attack on the Classical Binary 231

II. Why Some External Acts Deserve More Than the Protection Afforded to Expression 233

A. Speech Protection Exists for the Sake of the Thinker 233

B. Intellectual Privacy Protects Thinking, Reading, and Confiding 234

C. Surveillance Measurably Chills Inquiry 234

D. Thinking Happens Outside of the Skull 235

E. Four Features That Make an External Act Thought-Like 236

III. Testing the Analogies: Where Existing Doctrine Protects the Thinking Process, and Where It Fails 237

A. The Search Query: The Prompt's Ancestor 237

B. Confidential Dialogue: Jaffee and the Privilege Model 240

C. Reading: Stanley, Lamont, and the Right Not to Be Asked 241

D. Aggregated Records: Carpenter, Chatrie, and the CJEU on Metadata 241

E. The European Layer: Correspondence, Data Protection, and Cross-Border Production 243

F. Drafts and Private Writings: Provisionality Protected 245

G. Neural Data: Neurorights Principles Build an Intermediate Tier 247

H. The Convergence of Failures: Every Analogy Fails for a Similar Reason 249

IV. The U.S. Courtroom and the European Union Rulebook 251

A. 20 Million Conversations Under Ordinary Balancing 251

B. Heppner Against Warner: The Privilege Debate Asks the Wrong Question 252

C. Thought Moderation: Filtering as Interference with Formation 255

V. The Forum Interstitiale: Definition, Boundaries, and Protection Standard 259

A. Definition and Scope: The Anteroom Between the Fora 259
B. The Protection Standard: Five Elements 260
C. The Dilution Objection: A Qualified Tier Protects the Absolute Core 262
VI. Mapping Individual AI Actions Across the Three Fora 263
VII. Objections and Answers 265
A. Why Not Simply Repair the Binary? 265
B. Line-Drawing: Borrowed Boundaries, Not New Ones 266
C. The Relational Objection: Privileges Protect Functions, Not Relationships 266
D. Does Everything Fall Under Freedom of Thought? The Exclusion Criteria 267
E. The Seriousness Threshold: Should Mundane Prompts Even Qualify? 267
F. The Downgrade Risk and the No-Demotion Rule 268
Conclusion: The Need to Fill the Gap 268

INTRODUCTION: WHEN THREE U.S. COURT AI DECISIONS CREATE UNCERTAINTY

Between January and June of 2026, three courts in the United States reached conclusions that cannot comfortably coexist. On January 5, Judge Sidney Stein issued an order compelling OpenAI to produce 20 million de-identified consumer ChatGPT conversation logs to plaintiffs in a copyright case, upholding the magistrate judge's balancing of the privacy interests of users against the relevance of the documents.[1] On February 10 and 17, two federal district courts took opposite positions on whether a litigant's exchanges with a generative artificial intelligence (AI) system deserve protection from compelled disclosure: The Eastern District of Michigan shielded a self-represented plaintiff's ChatGPT queries as work product

[1] *In re* OpenAI, Inc., Copyright Infringement Litig., No. 25-md-03143 (S.D.N.Y. Jan. 5, 2026) (affirming the orders of Magistrate Judge Ona T. Wang of November 2025); *see* Kyle Jahner, *OpenAI Must Turn Over 20 Million ChatGPT Logs, Judge Affirms*, BLOOMBERG L. (Jan. 5, 2026, at 18:40 ET), https://perma.cc/F2V9-STPM. The plaintiffs had initially sought a sample of 120 million logs; the 20-million-log sample represents approximately 0.5 percent of the logs preserved by OpenAI in the litigation.

because it reasoned that generative AI programs are tools, not persons,[2] while the Southern District of New York ordered a criminal defendant to hand over 31 documents he had generated with Anthropic's Claude chatbot because the court determined that the defendant had no reasonable expectation of confidentiality for his exchanges with Claude due to Anthropic's terms of service and because Claude is not an attorney.[3] Then, on June 29, the Supreme Court held in *Chatrie v. United States* that the acquisition of Google Location History through a geofence warrant constitutes a Fourth Amendment search, because individuals retain a reasonable expectation of privacy in records of their movements even when a third party holds those records in the cloud.[4]

Read together, these decisions exemplify a structural gap. The most intimate records a person now generates—the questions they ask, the hypotheses they test, the drafts they abandon, and the fears they confide to a conversational system at three in the morning—received less protection in early 2026 than a log of where their phone had been. In other words, over the same six months, the movements of the body acquired constitutional protection in the United States while the movements of the mind were being handed over by the millions.

In parallel to the discussion before the courts, Sam Altman, CEO and cofounder of OpenAI, confronted with the preservation and production orders in the OpenAI litigation, called for the creation of an AI privilege based on his belief that talking to an AI "should be like talking to a lawyer or a doctor."[5] Altman's call has quite rightly been met with the objection that privileges protect relationships between humans and professionals who bear enforceable duties, and that an AI chatbot is neither human nor bears such duties. I treat this line of argumentation as a diagnostic rather than as a debate to be won: Both positions assume that protection for the human-AI conversation must be derived from the machine's status, and that assumption is, in my view, where the debate goes wrong. Part IV returns to this exchange in more detail. For now, the point is that the protective interest at stake

---

[2] Warner v. Gilbarco, Inc., No. 2:24-cv-12333, 2026 WL 373043 (E.D. Mich. Feb. 10, 2026).

[3] United States v. Heppner, No. 25-cr-00503-JSR, 2026 WL 436479 (S.D.N.Y. Feb. 17, 2026).

[4] 146 S. Ct. 2193 (2026).

[5] Sam Altman (@sama), X (June 5, 2025, at 20:33 ET), https://perma.cc/2UCP-Q26X ("[In my opinion,] talking to an AI should be like talking to a lawyer or a doctor."). The proposal and the responses to it are examined *infra* Part IV.B.

belongs to the thinking human, not to the machine or its operator, and that it calls for a category grounded in freedom of thought rather than for a privilege grounded in personification of the AI.

These differences in approach and the diverging results they generate stem from a legal and doctrinal architecture in international human rights that only offers two categories: on the one hand, freedom of thought, which benefits from absolute protection, and on the other, freedom of expression that can be subject to limitations. The current doctrine maps the two rights onto two spaces: the forum internum, described as the inner realm of thought, conscience, and belief; and the forum externum, which covers the outer realm of manifestation and expression.[6] Any cognitive act that leaves the skull falls, by default, into the outer category or out of the framework entirely into a residual class—that of business records held by a third party. Under this framework, a prompt typed into a chatbot is generally not considered to be inner thought, so it is not granted the absolute protection of the forum internum. As it resembles expression, some will consider it deserving of the qualified protection of speech. Others, as the OpenAI production order shows, consider it to qualify as a mere commercial record, which, in practice, means that it benefits from little protection.

This Article identifies the need to create and define a third category, the forum interstitiale, a distinct zone of heightened but qualified protection for cognition-adjacent acts—namely, external and technologically mediated behaviors that perform the function of thinking rather than the function of communicating. The zone sits between the two classical fora and borrows its structure from doctrines that already exist in fragmented form: the treatment of search queries as intellectual privacy, the constitutional protection of reading, the evidentiary privileges that attach to confidential dialogue, the recognition that communications metadata can reveal the content of a mind, and the emerging instruments that protect neural data. This Article's contribution is to name this gray zone, specify its boundaries, define the necessary protection standard, and map the range of individual AI interactions across the three fora.

---

[6] International Covenant on Civil and Political Rights art. 18, Dec. 16, 1966, 999 U.N.T.S. 171; Convention for the Protection of Human Rights and Fundamental Freedoms art. 9, Nov. 4, 1950, 213 U.N.T.S. 221; Charter of Fundamental Rights of the European Union art. 10, 2012 O.J. (C 326) 391.

Two design choices distinguish this proposal from adjacent positions in the literature. First, the forum interstitiale is qualified rather than absolute. The most forceful objection to expanding freedom of thought, pressed by Jan Christoph Bublitz and others, is that widening the absolutely protected sphere dilutes and jeopardizes the absoluteness of its core.[7] The interstitial zone concedes this point rather than fighting it, by making room for cognition-adjacent acts to receive heightened protection subject to narrow, predictable, and proportionate exceptions, while the traditional forum internum remains untouched and absolute. Second, the zone operates on a no-demotion rule. It lifts acts up from the forum externum or from unprotected third-party status; it does not create a mechanism that pulls absolutely protected acts down from the forum internum. Without that guardrail, an intermediate tier could become an instrument of downgrade in the hands of a court or a legislator, which would be antithetical to this proposal's intention.

The argument proceeds as follows. Part I sets out the classical binary and the scholarship that has established it. Part II supplies the theoretical case for treating some external acts as thought-like, drawing on thinker-based speech theory, intellectual privacy, empirical research on chilling effects, and the extended mind thesis, according to which cognitive processes can be partly constituted by resources outside of the head. Part III tests the candidate analogies in existing law, asking of each where it protects the thinking process and where it fails, and shows that the failures converge. Part IV turns to human-AI interaction itself, U.S. courtrooms where interstitial acts are being forced into ill-fitting categories, and the European rulebook where their future treatment is being written. Part V defines the forum interstitiale, its boundaries, and its protection standard, and answers the dilution objection. Part VI presents a taxonomy of individual AI actions, and Part VII confronts the remaining objections.

A note on method, scope, and contribution: The method employed herein is doctrinal and comparative. It starts from what legal texts and courts already say, and it moves in three steps. Part II.E draws four features from the theoretical liter-

---

[7] *See* Jan Christoph Bublitz, *Freedom of Thought in the Age of Neuroscience: A Plea and a Proposal for the Renaissance of a Forgotten Fundamental Right*, 100 ARCHIV FÜR RECHTS- UND SOZIALPHILOSOPHIE 1 (2014); Sjors Ligthart et al., *Rethinking the Right to Freedom of Thought: A Multidisciplinary Analysis*, 22 HUM. RTS. L. REV. 1 (2022).

ature that mark an external act as part of the thinking process: The act is provisional, it is not meant to persuade or inform, no human audience receives it, and it serves as an act of exploration. Part III holds the closest analogies in existing law against those features: the search query, confidential dialogue, protected reading, aggregated records, the European law on correspondence and data protection, the private draft, and neural data. None of these doctrines was written for the AI chatbot prompt, but each governs one of its nearest relatives. By looking at them, however, the Article comes to the conclusion that each fails in the same way for the purpose of the interstitial zone because the protection available to them is linked to a feature incidental to the thinking, never to the thinking itself. Part V then assembles the correct protection standard for the forum interstitiale from elements of those doctrines that have already been tested in court.

The status of users' exchanges with AI systems exists in multiple legal systems. It spans the International Covenant on Civil and Political Rights (ICCPR), the European Convention on Human Rights (ECHR), and the Charter of Fundamental Rights of the European Union (CFREU), and draws on national constitutional laws, mainly from the U.S. and Germany. The threat to our freedom to think when using AI is also two-directional. In one direction, AI systems can influence thought formation through flattery, engineering dependence on their assistance, and unseen steering of their users by way of refusing prompts or filtering responses. In the other direction, legislation and regulations can affect thought formation when liability pushes AI model providers[8] to filter and refuse what a person may explore with a machine.[9] This leads to the note on contribution: In my wider research, I call the first direction "thought manipulation" and the second "thought moderation" (a term deliberately chosen as the sibling of "content moderation"). I illustrate both directions for a general audience as four concrete mechanisms in the companion

[8] Throughout this Article, "AI model provider" means the entity that determines how a conversational AI system handles user interactions, whether by developing the model, deploying the model under its own terms, or both. The term is used in that ordinary sense and not in the narrower senses it carries in the Artificial Intelligence Act or the e-Evidence Regulation, both of which define the term for their own purposes.

[9] Caroline De Cock, *From Upload to Output Filters? How the AI Act's Code of Practice Could Threaten Freedom of Expression and Thought*, TECH POL'Y PRESS (Jan. 17, 2025), https://perma.cc/D2ZK-8M3F.

project to this Article: the assistant that flatters, the filter that refuses, the autocomplete that steers, and the answer box that replaces a debate with a paragraph.[10] The forum interstitiale is intended to cover the risks posed by both directions, because a protective category that guards only one flank misses half the problem.

I. The Classical Binary and Its Critics: Two Fora and a Gap in the Middle

A. The Architecture of the Two Fora

Article 18(1) of the ICCPR guarantees freedom of thought, conscience, and religion, including the freedom for a person to manifest their religion or belief. Article 18(3) permits limitations, but only to the freedom to manifest. This structure implies the distinction that scholarship and jurisprudence have made explicit: an inner freedom that admits no limitation, and an outer freedom that admits proportionate restriction.[11] The inner freedom is non-derogable under Article 4(2), even in public emergency. The Human Rights Committee confirmed that reading in General Comment No. 22, which states that Article 18 does not permit any limitations whatsoever on the freedom of thought and conscience, nor on the freedom to have or adopt a religion or belief of one's choice.[12] ECHR Article 9 and CFREU Article 10 replicate this architecture, and the Council of Europe's guidance describes the forum internum protected by Article 9 as absolute, meaning that the inner, private realm of an individual's mind—including their freedom of thought, conscience, religion, and personal belief—enjoys unqualified, non-derogable protection under international human rights law.[13]

---

[10] The two directions and the four mechanisms are developed for a general audience in the companion project to this Article, Caroline De Cock, *Free to Think: Terms and Conditions Apply*, FORUMINTERNUM.ORG, https://perma.cc/TU32-6E2N. The nearest prior term of art is the "cognitive content moderation" theorized in Mason Marks, *Cognitive Content Moderation*, 76 FLA. L. REV. 469 (2024). "Thought moderation" as it is used here is the broader term, reaching the pre-emptive restriction of private generative outputs.

[11] On the structure of Article 18 and the absolute character of its internal dimension, see Sjors Ligthart, *Reconsidering the Absolute Nature of the Right to Freedom of Thought*, 25 HUM. RTS. L. REV. 1 (2025).

[12] U.N. Hum. Rts. Comm., *General Comment No. 22: Article 18 (Freedom of Thought, Conscience or Religion)*, ¶ 3, U.N. Doc. CCPR/C/21/Rev.1/Add.4 (Sept. 27, 1993), https://perma.cc/G38Y-2SWH.

[13] EUR. CT. H.R., GUIDE ON ARTICLE 9 OF THE EUROPEAN CONVENTION ON HUMAN RIGHTS (Mar. 5, 2026), https://perma.cc/WJ7S-WCM9.

The first substantive United Nations treatment of freedom of thought since 1948 came in 2021, when Special Rapporteur Ahmed Shaheed presented his report on the right to the General Assembly.[14] The report identifies four attributes of the forum internum protection: (1) the freedom not to disclose one's thoughts; (2) the freedom from punishment for one's thoughts; (3) the freedom from impermissible alteration of one's thoughts; and (4) a positive obligation on states to foster an enabling environment for thought. The first and third attributes bear directly on this Article's subject. A regime in which conversational AI logs are able to be ordered into discovery for a court case on the basis of ordinary relevance balancing is, functionally, a regime of compelled disclosure of the record of thinking; a regime in which model outputs are pre-emptively filtered for compliance is a regime that shapes what may be thought with what will possibly become the most widely used thinking tools ever built. And in both cases, the fourth attribute does not seem fulfilled either.

The architecture of the forum internum rested, from the beginning, on a factual premise as much as a normative one. William Blackstone wrote that no temporal tribunal "can search the heart, or fathom the intentions of the mind" and therefore "cannot punish for what it cannot know."[15] Two centuries later, the U.S. Supreme Court could still describe the mind and spirit of man as "forever free," beyond the reach of judgments "ineffective to direct the course of man."[16] The forum internum has historically been deemed absolute partly because it was thought unreachable and so protected, in effect, by its own opacity. It was considered a black box. Every technology examined in this Article strains the limits of that premise, and the reactions from courts and legislators run the risk of eroding our most fundamental freedom.[17]

---

[14] Ahmed Shaheed (Special Rapporteur on Freedom of Religion or Belief), *Interim Rep.: Freedom of Thought*, U.N. Doc. A/76/380 (Oct. 5, 2021), https://perma.cc/AA2N-F6BK. This formulation is standard in the doctrinal literature. *See* SUSIE ALEGRE, CTR. FOR INT'L GOVERNANCE INNOVATION, PROTECTING FREEDOM OF THOUGHT IN THE DIGITAL AGE (2021), https://perma.cc/Q588-KUZJ; Evelyn Mary Aswad, *Losing the Freedom to Be Human*, 52 COLUM. HUM. RTS. L. REV. 306 (2020).

[15] 4 WILLIAM BLACKSTONE, COMMENTARIES *21.

[16] Jones v. Opelika, 316 U.S. 584, 593, 595 (1942).

[17] JULIE E. COHEN, *The Future(s) of Fundamental Rights*, *in* BETWEEN TRUTH AND POWER: THE LEGAL CONSTRUCTIONS OF INFORMATIONAL CAPITALISM 238 (2019).

B. Three Lines of Attack on the Classical Binary

The classical binary has come under sustained academic attack, from at least three fronts.

The first front is conceptual. In the most direct attack, Simon McCarthy-Jones and Joel Walmsley argue that thought does not occur only in a forum internum inside our heads, and that the right to freedom of thought should recognize and protect two types of what they call "forum externum thought."[18] Type 1 is "extended thought," which occurs when we extend our mind using features of our body or environment, such as writing in a diary and searching on the internet. Type 2 is "thoughtspeech," which occurs when we speak with others to think aloud together. McCarthy-Jones and Walmsley propose that these forms of external thought should receive absolute protection when employed for the purposes of self-government and truth-seeking, using criteria of trust and functional integration drawn from extended mind theory to mark the boundary. This theory matters for our purposes because it names, from within the freedom of thought literature itself, acts that are very close to what a person performs with a conversational AI system: querying and thinking aloud in dialogue. My Article departs from McCarthy-Jones and Walmsley on the protection standard that should be applied to these behaviors. Their proposal extends absolute protection into the forum externum, conditioned on the purposes of self-government or truth-seeking. That condition is designed to rebut the objection, pressed by Bublitz and others, that widening the absolute sphere dilutes it. A proposal that extends absoluteness outside the forum internum must still defend the absoluteness of the acts it seeks to cover via the extension. In contrast, this Article's formulation of the forum interstitiale reaches the same acts through a qualified tier. Part V.C returns to this point.

The second front is doctrinal. Sjors Ligthart has challenged the assumption that the right to freedom of thought is completely absolute, proposing a narrow set of implied limitations on certain elements of the right, provided that they do not impair its very essence.[19] Caroline Roberts, on a close reading of the European Court of Human Rights case law under ECHR Article 9, concludes that the inner sanctum

---

[18] Simon McCarthy-Jones & Joel Walmsley, *What Is Thought and What Makes It Free? Or, How I Learnt to Stop Worrying and Love the* Forum Externum, *in* THE CAMBRIDGE HANDBOOK OF THE RIGHT TO FREEDOM OF THOUGHT 305 (Patrick O'Callaghan & Bethany Shiner eds., 2025).

[19] Ligthart, *supra* note 11.

supposedly receiving hands-off treatment is simply not evident in the Court's practice, and describes the jurisprudence instead as loose concentric circles: an innermost circle of strong forum internum relevance that receives the highest protection; an outermost circle that receives very little; and a contested middle circle in which forum internum relevance and countervailing state interests are both strong, and in which the Court attempts to reach a balance between those two forces.[20] Roberts's middle circle is the closest existing doctrinal precursor to the forum interstitiale: a descriptive finding that an intermediate zone already operates unnamed inside the case law of the very court most expected to adopt the absolutist reading.

The third front is cautionary. Bublitz and some of his peers have warned that the distinction of whether a thought has been manifested externally or remains unmanifested is vulnerable to critique building on the theory of extended cognition, while also insisting that the response cannot be an indefinite expansion of the absolute sphere because a right that covers everything absolutely will end up protecting nothing absolutely.[21] This is the dilution objection, and it is the strongest argument in the field against protecting external cognitive acts as thought. The architecture proposed here is built to answer it, and Part V addresses it directly.

The critique of the boundaries between manifested and unmanifested thought has now reached the scholarship on freedom of religion and belief, where the forum internum doctrine was built in the first place. Writing in 2026, Andrea Pin argues that the assumption underpinning the whole architecture—that the mind is shielded by its own inaccessibility—cannot survive systems able to infer what a person thinks and feels, and that regulators preoccupied with prohibiting manipulation have addressed influencing techniques while leaving their object, *i.e.*, thought and conscience themselves, without direct protection.[22] Pin's analysis calls for a reappraisal of the current regulatory approach but leaves open what that reappraisal should produce, which is the gap this Article sets out to fill.

The current state of the relevant literature can therefore be summarized briefly. The binary between internality and externality is doctrinally prevalent, even though it is increasingly evident that such a black-and-white ideation of thought formation

---

[20] CAROLINE K. ROBERTS, FREEDOM OF RELIGION OR BELIEF IN THE EUROPEAN CONVENTION ON HUMAN RIGHTS (2023).

[21] Ligthart et al., *supra* note 7.

[22] Andrea Pin, *Freedom of Thought, Conscience, and AI*, 40 J.L. & RELIGION 454 (2026).

is ill-adapted to the realities of technological evolution. Its critics agree that acts outside the skull can constitute or condition thought, and disagree on what follows: absolute protection for a functionally defined subset of acts (McCarthy-Jones and Walmsley), implied limitations inside a nominally absolute right (Ligthart), honest recognition of the judicial balancing happening (Roberts), or restraint from full protection in the name of preserving the absolute core (Bublitz). A different response is available, however: one that accepts the binary and looks for a sharper test within it, one that I considered in Part VII. No participant in the debate has yet proposed what the fragmentary agreement within it points toward: a named, structured, intermediate category with its own boundaries and its own protection standard. That is the gap this Article sets out to fill by proposing the forum interstitiale as a novel legal category.[23]

II. WHY SOME EXTERNAL ACTS DESERVE MORE THAN THE PROTECTION AFFORDED TO EXPRESSION

If the interstitial zone is to emerge as a new category worthy of protection, it needs a theory of why certain external acts deserve a status above expression. Four bodies of work offer helpful elements for building such a theory.

A. Speech Protection Exists for the Sake of the Thinker

Seana Shiffrin's thinker-based approach to freedom of speech argues that the primary function of expressive protection is to safeguard the development and operation of the mind of the thinker, rather than the instrumental values of democratic participation or the marketplace of ideas.[24] On this view, speech matters because it is the central mechanism through which a person externalizes a thought and, in doing so, tests it, revises it, and comes to know her own mind. The two accounts Shiffrin sets aside both assume an audience, protecting speech because someone receives it, and hers does not. This is the normative step my argument

---

[23] The term does not, to my knowledge, appear in prior legal scholarship. The nearest structural precedents, the contested middle circle Roberts identifies descriptively in the Strasbourg case law and the core and periphery model Blitz proposes for U.S. constitutional rights, are discussed *supra* Part I.B and *infra* Part II.D, respectively; neither names, generalizes, or operationalizes the intermediate zone as a category of its own.

[24] Seana Valentine Shiffrin, *A Thinker-Based Approach to Freedom of Speech*, 27 CONST. COMMENT. 283 (2011); SEANA VALENTINE SHIFFRIN, SPEECH MATTERS: ON LYING, MORALITY, AND THE LAW (2014).

needs, because it gives a reason for protection to reach external acts that does not run through the presence of a listener.

B. Intellectual Privacy Protects Thinking, Reading, and Confiding

Neil Richards's work on intellectual privacy translates the thinker-based insight into an institutional claim: Freedom of thought requires a protected zone for the pre-expressive activities of thinking, reading, and confiding, because ideas are formed there before they are ever expressed.[25] Surveillance of intellectual records, search histories, reading habits, and communications is, on this account, especially harmful, and access to such records should require a showing of demonstrated need well above the ordinary standard. Richards wrote before conversational AI existed at scale, but his three activities map onto it with almost no translation: A human prompting a chatbot constitutes thinking, absorbing the text generated in response is reading, and engaging in a sustained personal exchange with the chatbot is confiding. The intellectual privacy literature is, in substance, a theory of the forum interstitiale developed within First Amendment scholarship a decade before the acts it so closely describes became mainstream.

C. Surveillance Measurably Chills Inquiry

The claim that surveillance of intellectual activity deters thinking is testable, and it has been tested. Jonathon Penney's study of Wikipedia traffic before and after Edward Snowden's June 2013 disclosures of mass surveillance by the U.S. National Security Agency found a statistically significant, immediate, and sustained decline in views of articles on terrorism-related topics, the primary categories a reasonable person would imagine likely to trigger government attention.[26] People did not merely speak less; they inquired less, which locates the injury to their freedom at the point where thought is formed, not where it is expressed. This finding matters doctrinally because it grounds the harm of potential intellectual surveillance at the formation stage of thought rather than at the expression stage, which is exactly where the classical binary of internality versus externality provides little to no protection. Every argument in this Article for shielding search queries and chatbot

---

[25] Neil M. Richards, *Intellectual Privacy*, 87 TEX. L. REV. 387 (2008); NEIL M. RICHARDS, INTELLECTUAL PRIVACY: RETHINKING CIVIL LIBERTIES IN THE DIGITAL AGE (2015).

[26] Jonathon W. Penney, *Chilling Effects: Online Surveillance and Wikipedia Use*, 31 BERKELEY TECH. L.J. 117 (2016); Jonathon W. Penney, *Understanding Chilling Effects*, 106 MINN. L. REV. 1451 (2022).

prompts rests empirically on the demonstrated fact that a lack of protection yields a reluctance to ask questions, resulting in a chilling effect on thought itself.

D. Thinking Happens Outside of the Skull

Andy Clark and David Chalmers's extended mind thesis holds that cognitive processes can be partly constituted by resources outside of the skull when those resources are functionally integrated into thinking, reliably available, and trusted.[27] That theory still treats thinking as something a mind does individually, albeit sometimes with help from the world around it. A second tradition drops even that assumption. Edwin Hutchins studied how ships are navigated and aircraft flown, and found that the thinking involved is not being done in any one person's head. It is done by a team working with charts and instruments, each part of the system holding a piece of the problem and passing it on. The title of one of his papers makes the point: "How a Cockpit Remembers Its Speeds."[28] This fits with the human-AI exchange better than the theory of thought extension does. In an iterative dialogue with an AI model, an idea is put out by the human user, altered by the machine in response, and returned for another pass. The thinking is done through this loop of activity rather than by either side alone. The loop is the fullest case of what this Article will refer to as the thinking process, which also includes acts with no feedback loop at all, such as the single query, the silent reading, the note that expects no answer. In every case, the undertaken behavior generates the record of a process rather than the announcement of a result. This framework also carries a risk that my thesis will attempt to disarm: If cognition is distributed across the loop of the thinking process, a critic might locate part of it in the AI model provider's model and derive a provider stake in the record of the thinking activity. The answer to this is that rights do not follow the mechanics used for processing, but rather that they follow the person. Distributed cognition describes how the thinking happens, while freedom of thought determines whose thinking it is; and the initiative, the purposes, and the vulnerability in the human-AI loop all belong to the human participant. This description of the thinking process strengthens the case for interstitial protection, because a record of the loop is a record of thinking in progress.

---

[27] Andy Clark & David J. Chalmers, *The Extended Mind*, 58 ANALYSIS 7 (1998).

[28] EDWIN HUTCHINS, COGNITION IN THE WILD (1995); Edwin Hutchins, *How a Cockpit Remembers Its Speeds*, 19 COGNITIVE SCI. 265 (1995).

Marc Blitz has spent two decades translating Hutchins's thesis into constitutional structure, arguing that devices integrated into cognitive routines deserve treatment closer to freedom of thought than to ordinary property or privacy law[29] and, crucially for present purposes, proposing that freedom of thought itself should be modeled with a core in which protection is close to absolute and a periphery in which the state retains narrow, structured flexibility.[30] Blitz's core and periphery supply the architectural precedent for a graduated right; Caroline Roberts's concentric circles supply the descriptive evidence that courts already administer one.[31] With the concept of forum interstitiale, this Article formalizes the structure suggested in both authors' work.

E. Four Features That Make an External Act Thought-Like

Looking at the existing gray zone between the protection afforded to the forum internum and the forum externum, an external act within the thought process seems to warrant more than a freedom of expression standard of treatment when it exhibits the functional profile of thinking.

Four features recur across the theoretical accounts and the doctrines examined in the next Part. The first feature is provisionality: The act is revisable and nonfinal, a step in formation rather than a finished position, in the way a draft or an exploratory question is provisional and a publication is not. The second is noncommunicative intent: The act is not directed at persuading or informing an audience, and its addressee, where one exists, functions as an instrument of the person's own thinking. The third is the absence of an audience: No human public receives the act, and the person reasonably expects that none will. The fourth is epistemic exploration: The act tests an idea rather than asserting one, including an idea the person may not hold and a position the person is trying on in order to make a decision about it, be it to embrace or reject it. An act bearing this profile is evidence of a mind at work rather than a message sent, and a legal system that treats the two

---

[29] Marc Jonathan Blitz, *Freedom of Thought for the Extended Mind: Cognitive Enhancement and the Constitution*, 2010 WIS. L. REV. 1049.

[30] Marc Jonathan Blitz, *Freedom of Thought and the Structure of American Constitutional Rights*, *in* 1 THE LAW AND ETHICS OF FREEDOM OF THOUGHT 103 (Marc Jonathan Blitz & Jan Christoph Bublitz eds., 2021).

[31] *See supra* note 23 and accompanying text.

identically will likely end up creating a chilling effect on thought in the name of regulating speech.

These four features also supply the test for the next Part. Each candidate doctrine is examined for how far its protection tracks these features; where it tracks something else instead—such as a relationship, a recording, an addressee, or a data type—the analogy fails.

III. Testing the Analogies: Where Existing Doctrine Protects the Thinking Process, and Where It Fails

The forum interstitiale does not require the invention of an entirely new legal construction. Legal systems already grant several classes of cognition-adjacent acts a protection at a level above ordinary expression or ordinary records, without ever naming the category they are building. The behaviors needing classification can be identified in light of current online and offline actions that use technology: information-seeking prompting, exploratory and hypothetical prompting, silent reading of generated output, iterative drafting and brainstorming, journaling, venting and therapy-like conversation, the accumulation of interaction memory, and, at the boundary of the interstitial zone, delegation to agents and the process leading to publication. For each, the lawyer's instinct is to look at existing analogies, examined below: the search query, confidential dialogue, protected reading, aggregation doctrine, the European law on correspondence and data protection, the private draft, and neural data. Each doctrine is tested against the four features of the functional profile of thinking laid out in Part II.E above. The analysis begins with the closest behavioral analogy, the search query, and moves outward.

A. The Search Query: The Prompt's Ancestor

The closest existing analogue in the law to interacting with a chatbot is entering a query into a search engine, and the legal status of such queries is therefore the natural starting point for exploring the legal regime available to AI-human interaction. This is becoming all the more true as search engines switch to AI-powered modes and, conversely, as people begin to use chatbots as search engines. The two acts involve the same initial behavior: A person types a question in natural language into a text box operated by a commercial provider, transmits it, and receives an answer. What the chatbot contributes—namely, the sharing of richer questions, a dialogue rather than a one-shot exchange, a generated answer rather than a ranked list, and persistent memory of past sessions—intensifies every element that made the query legally sensitive but does not fundamentally change any of them in kind.

Whatever protection the law settles on for the search query is accordingly the floor from which the protection for a prompt to a chatbot must be argued.

That legal protection floor is, at present, contested on every continent that has considered it. In the U.S., the doctrinal baseline is the third-party doctrine: Information voluntarily conveyed to a search provider carries no reasonable expectation of privacy. Two state supreme courts have now applied that baseline to the same investigative instrument—the reverse keyword warrant, which compels a search provider to identify every user who entered a given term—and reached opposite conclusions. In *People v. Seymour*, the Colorado Supreme Court held that the defendant had a constitutionally protected privacy interest in his Google search history even when revealed only in connection with his Internet Protocol address, that he retained a possessory interest in his search terms, and that because search history implicates freedom of expression, constitutional protections must be applied with scrupulous exactitude.[32] In doing so, the court built on its earlier holding that compelled disclosure of bookstore purchase records requires heightened justification.[33] Its move in *Seymour* was, in substance, interstitial: It treated queries neither as ordinary business records nor as public speech but as a third thing that demands a stricter standard because of what it reveals about the searching mind. Two years later, in *Commonwealth v. Kurtz*, a three-justice plurality of the Pennsylvania Supreme Court held the opposite view: Entering a query into a search engine voluntarily turns its contents over to a third party and negates any constitutionally recognized privacy interest and the need for a warrant.[34] Three other justices concurred in judgment only, determining that Pennsylvania statutorily mandated a warrant for such queries but that probable cause for a warrant had been met.[35] The split is the classification problem in miniature. The same act, examined through the same instrument, is qualified as protected intellectual activity in Denver and potentially as an unprotected business record in Harrisburg because the law currently possesses no category that explicitly says what the legal standing of a search query is. Beneath the legalistic layer, practice is thankfully more protective: Search providers generally treat queries as content demanding a warrant under the federal

---

[32] People v. Seymour, 536 P.3d 1260 (Colo. 2023) (quoting Zurcher v. Stanford Daily, 436 U.S. 547, 564 (1978)).

[33] Tattered Cover, Inc. v. City of Thornton, 44 P.3d 1044 (Colo. 2002).

[34] Commonwealth v. Kurtz, 348 A.3d 133 (Pa. 2025).

[35] *Id.* at 156.

stored communications regime, and California requires one by statute.[36] Search histories are nonetheless routinely introduced at trial as evidence of intent and premeditation, the practical manifestation of the thoughtcrime concern: that people will be punished for a thought they had even though they never acted on it.

Outside of the U.S., Canada has moved furthest in the protective direction. The Supreme Court of Canada rejected the third party logic in *R. v. Spencer*, protecting the subscriber information that lies behind an Internet Protocol address precisely because its disclosure unmasks intimate online activity, and thus supporting a right to online anonymity.[37] It then extended constitutional protection to IP addresses themselves in *R. v. Bykovets*.[38] Under the Canadian approach, which asks whether information reaches the biographical core of the person, the search query would likely sit squarely in the protected zone.

Europe, by contrast, protects the search query through layered regimes rather than via a dedicated doctrine. Queries are personal data, and queries revealing a person's belief, health, or sex life fall within the special categories of Article 9 of the General Data Protection Regulation (GDPR). When transiting over a communications network, they are confidential communications within the ePrivacy regime; and the metadata jurisprudence examined below treats usage records as capable of being no less sensitive than content. The European Court of Human Rights has held that identifying the user behind a dynamic Internet Protocol address engages Article 8 of the ECHR, which protects private life and correspondence, precisely because the address unmasks what was searched and read by its user.[39] The German Federal Constitutional Court has gone furthest in terms of protection for the search query. In 2008, it derived a distinct fundamental right to the confidentiality and integrity of information technology systems. This is the closest any legal order has come to recognizing an intermediate category of the kind this Article proposes.[40]

The search query, in short, is everywhere sensed to be more than a record and nowhere given a name that says what that "more" is. Doctrinally, it is claimed by

---

[36] California Electronic Communications Privacy Act, CAL. PENAL CODE § 1546 (West 2016).

[37] R. v. Spencer, [2014] 2 S.C.R. 212 (Can.).

[38] R. v. Bykovets, 2024 S.C.C. 6 (Can.).

[39] Benedik v. Slovenia, App. No. 62357/14 (Eur. Ct. H.R. Apr. 24, 2018), https://perma.cc/4U8M-Z8KS.

[40] BVerfG, 1 BvR 370/07, 1 BvR 595/07, Feb. 27, 2008, https://perma.cc/N64U-BD83.

expression in Denver, by ordinary records according to a judicial plurality in Harrisburg, by anonymity in Ottawa, by data protection in Brussels, and by a *sui generis* computing right in Karlsruhe. The chatbot prompt inherits this entire unsettled estate with every sensitivity in how we engage amplified: The prompt is often longer and more self-revealing than the search query, the dialogue generally accumulates where the query stood alone, and the chatbot's generated answer feeds back into the user's next thought, such that the log records not a moment of curiosity but a thinking process.

Thus, the search query supplies the floor of a legal standard for the interstitial zone but cannot supply the ceiling. It works as an analogy to the chatbot prompt because the conduct of the two actions is similar, it brings with it the intermediate standards that courts have already built, and its case law shows that records revealing how a mind formed a view can attract heightened protection. It fails, however, on two counts. It inherits the voluntary disclosure argument that partially succeeded in the Pennsylvania case; and it describes a single act, which leaves it silent on everything the prompt adds: the dialogue between human and chatbot, the memory of those exchanges, and the user's confessional register.

B. Confidential Dialogue: Jaffee *and the Privilege Model*

For the confessional register of conversational AI, the governing analogy is the law of evidentiary privilege for confidential dialogue. In *Jaffee v. Redmond*, the U.S. Supreme Court recognized a federal psychotherapist-patient privilege, reasoning that effective psychotherapy depends upon an atmosphere of confidence and trust in which the patient is willing to make a frank and complete disclosure of facts, emotions, memories, and fears, and that the mere possibility of disclosure may impede the development of the relationship that this privilege exists to protect.[41] The Court rejected case-by-case balancing on the ground that an uncertain privilege is little better than no privilege at all, and weighed the evidential cost against public mental health, which it called a public good of transcendent importance. Attorney-client and priest-penitent privileges follow the same logic against the background presumption that the public has a right to every man's evidence. This privilege model contributes two elements to the interstitial standard: Protection must be predictable ex ante if it is to do its work of enabling candor and removing any and all

[41] Jaffee v. Redmond, 518 U.S. 1, 10, 17–18 (1996).

chilling effects, and exceptions must be narrow and standardized rather than discretionary. Where the analogy fails is equally clear, and Part IV examines the failure in more detail: The privileges as they stand are tied to a human professional owing enforceable duties, which no consumer AI system is or owes, and privilege is merely an evidentiary rule, silent on retention, reuse, and filtering. In other words, the analogy supplies the shape of the protection, but it cannot supply its basis.

C. Reading: Stanley, Lamont, *and the Right Not to Be Asked*

The U.S. Supreme Court held in *Stanley v. Georgia* that the state has no business telling a man, sitting alone in his own house, what books he may read or what films he may watch, and that the Constitution protects the right to receive information and ideas.[42] *Lamont v. Postmaster General* struck down a requirement that recipients affirmatively request material the government had labeled communist political propaganda, protecting readers from having to identify themselves as such.[43] The library confidentiality statutes enacted in the U.S., and the reader privacy legislation that followed them, extend the same principle to records of reading.[44] The freedom to read approach establishes that consuming ideas is not expression and is not unprotected conduct either; it is an input to thought, and inquiries into it are policed at a level that ordinary conduct never attracts. This analogy holds fully for the silent reading of output generated by a chatbot, which is why it sits high in the taxonomy's possible analogies. It fails, however, for everything upstream of the reading: The person interacting with a model is not only a recipient but a co-producer of the text they read, and reader protection has no relevance for prompts given to a chatbot, which are the user's own contributions and can therefore be the most revealing part of the log.

D. Aggregated Records: Carpenter, Chatrie, *and the CJEU on Metadata*

The U.S. Supreme Court has now recognized that accumulated interaction records can reveal the content of a mind even when a single record does not. In *Carpenter v. United States*, it held that historical cell site location information is protected notwithstanding its custody by a third party, because it provides a detailed

---

[42] Stanley v. Georgia, 394 U.S. 557, 564–65 (1969).

[43] Lamont v. Postmaster General, 381 U.S. 301 (1965).

[44] *See, e.g.*, California Reader Privacy Act, CAL. CIV. CODE § 1798.90 (2011). On the library confidentiality statutes, see generally *State Privacy Laws Regarding Library Records*, AM. LIBR. ASS'N (Nov. 2021), https://perma.cc/LS4S-6GZF.

chronicle of a person's physical presence.[45] In June 2026, the Court extended that logic in *Chatrie*, holding that the acquisition of Google Location History through a geofence warrant is a Fourth Amendment search and rejecting the argument that brevity of the tracking period defeats the expectation of privacy.[46] In Europe, the Court of Justice of the European Union (CJEU) invalidated the EU Data Retention Directive via *Digital Rights Ireland* partly because retained metadata, taken as a whole, may allow very precise conclusions to be drawn concerning the private lives of the persons whose data has been retained, and partly because blanket retention is likely to lead the persons concerned to feel that their private lives are the subject of constant surveillance.[47] The CJEU added in *Tele2 Sverige* that metadata is no less sensitive, in terms of the right to privacy, than the actual content of communications.[48] In *La Quadrature du Net II*, the CJEU permitted access to Internet Protocol address data for copyright enforcement, only under strict conditions, including the partitioning of identity data from data that reveals what a person accessed online, and held that prior review of access to the collected data could in no case be fully automated.[49] These aggregation cases contribute a third element to the interstitial standard: Protection must be sensitive to mosaics. A single prompt may reveal little; a year of prompts is a mind in longitudinal section, and access rules must scale with what the aggregate discloses.

This analogy reaches its limit here because aggregation doctrine protects the pile and not the piece. Its threshold is quantitative, so a single exploratory prompt falls below it, no matter how revealing that prompt may be on its own. And the pile the location cases protect is only a chronicle of where a person's body went; extending the same protection to a record of what the person's mind asked is exactly what the Pennsylvania court declined to do in *Kurtz*. What the doctrine does supply,

---

[45] 585 U.S. 296 (2018).

[46] 146 S. Ct. 2193 (2026). The Court held that a search had occurred and remanded the questions of probable cause and particularity to the Fourth Circuit, so the decision settles classification—which is the point relevant here—while leaving warrant validity open.

[47] Joined Cases C-293/12 & C-594/12, Digit. Rts. Ir. Ltd. v. Minister for Commc'ns, ECLI:EU:C:2014:238, ¶¶ 27, 37 (Apr. 8, 2014), https://perma.cc/TQ8B-A3UH.

[48] Joined Cases C-203/15 & C-698/15, Tele2 Sverige AB v. Post-Och Telestyrelsen, ECLI:EU:C:2016:970, ¶ 99 (Dec. 21, 2016), https://perma.cc/73RN-E2B7.

[49] Case C-470/21, La Quadrature du Net v. Premier Ministre, ECLI:EU:C:2024:370 (Apr. 30, 2024), https://perma.cc/FU3S-XCHE.

though, is the aggregation caveat itself. A system that retains a year of a person's exploration holds exactly the kind of accumulation these cases were built for, which is why the taxonomy coming in Part VI marks AI memory and persistent interaction profiles as aggregation-sensitive, while a single prompt has to draw its protection from elsewhere in the interstitial standard.

E. The European Layer: Correspondence, Data Protection, and Cross-Border Production

Three further European instruments protect interstitial acts without naming them, and together they show that the intermediate tier is not just a U.S. artifact. The first is ECHR Article 8, which protects private life, home, and correspondence. The European Court of Human Rights has long read "correspondence" as following the forms communication actually takes, from workplace email and internet usage records to instant messaging, and its surveillance jurisprudence subjects the interception of communications to requirements of foreseeability, necessity, and independent oversight.[50] Whether a sustained exchange with a conversational system counts as correspondence under Article 8 is an open question. An affirmative reading would import the Strasbourg-based Court's surveillance case law into human-AI interaction and would go a long way toward developing the interstitial zone proposed here. The functional case for this move is strong because the European Court has consistently attached protection to the confidentiality of the channel rather than to the identity or status of the party at the other end, which is the same move that Part IV defends against the relational frame. An affirmative reading would not, however, impart all of the protection needed for the forum interstitiale. Article 8 protection is qualified, its balancing is ordinary, and it binds only Convention states. More importantly, the channel is still not the act. Protecting the confidentiality of an exchange says nothing about the second direction of the problem, as interference in the thinking loop of a human-AI interaction can come not only from an exposure of what the user has said but from a filtering of what the user can be told in response to their prompt.

The second European instrument is data protection. Under the GDPR, conversational logs are personal data, and a corpus of extended exchanges will routinely

---

[50] Copland v. United Kingdom, App. No. 62617/00 (Eur. Ct. H.R. Apr. 3, 2007), https://perma.cc/LKB8-NETC; Barbulescu v. Romania, App. No. 61496/08 (Eur. Ct. H.R. Sept. 5, 2017), https://perma.cc/BM4H-NNUU; Big Brother Watch v. United Kingdom, App. Nos. 58170/13, 62322/14 & 24960/15 (Eur. Ct. H.R. May 25, 2021), https://perma.cc/2ZFV-P84F.

reveal a person's religious and philosophical beliefs, political opinions, health, and sex life, all special categories to which GDPR Article 9 attaches a general prohibition on processing.[51] Purpose limitation constrains the reuse of conversations for ends that their author never contemplated, storage limitation constrains their retention, and the right to erasure gives the author a claim to their destruction. The GDPR is, in substance, an interstitial instrument: It grants heightened, structured, and qualified protection to records because of what they reveal, without asking whether they are thought or expression. It also produces the sharpest transatlantic collision in this field because the U.S. preservation order examined in Part IV compels an AI model provider to retain, for litigation use, the very conversations that European law obliges it to minimize the storage of and, on request, to delete.

The third is compelled production, where the European Union is about to operate its own pathway. The e-Evidence Regulation allows a judicial authority in a Member State to order any service provider offering services in the Union to produce or preserve electronic evidence, content data included, within ten days or, in emergencies, eight hours, with the regime applicable starting on August 18, 2026.[52] AI conversation logs are content data within its scheme. The practical implementation of the framework remains unclear, including whether orders under it treat conversational records as ordinary content or as records that require stronger justification before they can be compelled. The outcome will be an early, consequential test of the classification this Article proposes.

The European analogies carry their own failures. Correspondence presumes an addressee: The Strasbourg case law protects the channel between persons, and whether it stretches to an exchange in which one party is not a human has not yet been decided, so the EU law acquis will be built only if a court first accepts the analogical step. Conversely, data protection, for all its reach, protects the person as a data subject rather than their thought process itself: It qualifies nearly everything and forbids very little, its lawful bases and derogations run straight through the conversational corpus, and the CJEU has already permitted access to usage data for

---

[51] Regulation 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data (General Data Protection Regulation), arts. 5(1)(b), 5(1)(e), 9, 17, 2016 O.J. (L 119) 1, https://perma.cc/9E8A-TUCT.

[52] Regulation 2023/1543 of the European Parliament and of the Council of 12 July 2023 on European Production Orders and European Preservation Orders for Electronic Evidence in Criminal Proceedings, 2023 O.J. (L 191) 118, https://perma.cc/K5FU-RDVS (effective Aug. 18, 2026).

interests as ordinary as copyright enforcement. The GDPR supplies constraints on reuse and retention; it does not supply a categorical answer to what a conversational record is, and it certainly does not do much to protect such a record in court.

F. Drafts and Private Writings: Provisionality Protected

The U.S. Supreme Court held in the 1886 case *Boyd v. United States* that compelling a person to produce his private books and papers for use against him amounted to both a search and seizure and compelled self-incrimination, such that the Fourth and Fifth Amendments ran together to shield private papers from the state.[53] That broad rule did not survive, but the act of production doctrine that replaced it kept a narrower form of the idea: The compelled act of producing documents is itself testimonial, because it concedes that the papers exist, that the person holds them, and that they are what the subpoena says they are. What *Boyd*'s lineage bequeaths is therefore an intuition rather than a rule: that a person's private papers, whether a diary, a note, or a draft, are things the state should reach only with difficulty, because compelling their production conscripts the mind that made them. The work product doctrine performs a parallel function in civil litigation, shielding materials that reveal mental impressions and litigation strategy. Drafts are the paradigm of provisionality: writing that exists in order to be revised, that represents positions under consideration rather than adopted, and that has no reader for whom it is the finished communication. Circulation for comment does not by itself end that condition, since a colleague who reads a draft in order to improve it is a participant in the writing rather than its final destination. The sticking point in a work product claim is therefore not whether a human has read the draft but how far their reading of it has carried it toward a public.

The interstitial claim of the iterative AI-assisted draft thus aligns with a very old intuition about the privacy of unfinished writing. Wherever institutions deliberate before deciding, the law has built something to shield such work. The deliberative process privilege shields an agency's pre-decisional discussions so that officials may explore options candidly; judicial and jury deliberations are secret; cabinet confidentiality is a constitutional convention; and the European Union's access-to-doc-

---

[53] Boyd v. United States, 116 U.S. 616 (1886). On the modern act of production doctrine, *see* Fisher v. United States, 425 U.S. 391 (1976).

uments regime permits its institutions to protect their own decision-making process.[54] Each of these protects documents that are passed among several people, because passing them around is part of the deliberative process, rather than the end of it. Legal orders on both sides of the Atlantic long ago accepted that candid thinking requires a protected process and reserved that protection for the thinking of the state. My proposal extends the same reasoning to a person's thinking as an individual.

The individual has fared worse. The diary analogy has already failed, and, in the U.S., what survives of *Boyd* protects only the act of producing documents and not their contents. Even that act is privileged only where producing would tell the state something it does not already know; where the existence and authenticity of the papers are a foregone conclusion, the privilege falls away and the contents, never protected to begin with, are simply available to be read. A sitting senator's personal diaries were ordered to be produced on exactly that reasoning.[55] European Court of Human Rights case law reaches the same result by another route, holding that the privilege against self-incrimination does not extend to material with an existence that is independent of the alleged suspect's will, which is what every pre-existing diary is.[56] Germany litigated the boundary itself. The Federal Constitutional Court's core area of private life, which it grounds in the guarantee of human dignity under Article 1 of the Basic Law (Grundgesetz), allows no balancing at all, and the Court set out the privacy doctrine as applied to covert surveillance in 2004. In 2011, the German Federal Court of Justice applied this principle to thinking aloud to oneself, excluding a suspect's secretly recorded self-talk from evidence even in a murder case.[57] Yet in the *Tagebuch* decision of 1989, an evenly divided Court allowed a

---

[54] NLRB v. Sears, Roebuck & Co., 421 U.S. 132 (1975) (deliberative process privilege under Freedom of Information Act exemption 5); Regulation 1049/2001 of the European Parliament and of the Council of 30 May 2001 Regarding Public Access to European Parliament, Council and Commission Documents, art. 4(3), 2001 O.J. (L 145) 43, https://perma.cc/G32T-7UB9.

[55] *Fisher*, 425 U.S. at 410; United States v. Doe, 465 U.S. 605, 612 (1984); Senate Select Comm. on Ethics v. Packwood, 845 F. Supp. 17, 22–23 (D.D.C. 1994).

[56] Saunders v. United Kingdom, App. No. 19187/91 (Eur. Ct. H.R. Dec. 17, 1996), https://perma.cc/JA4G-W6J8.

[57] BVerfG, 109 BVerfGE 279, Mar. 3, 2004, https://perma.cc/D3F9-M3RJ (finding, in the context of acoustic surveillance of the home, that the core area of private life is grounded in Article 1 of the Basic Law and admits no balancing); BGH, 2 StR 509/10, 2012 NStZ 277, Dec. 22, 2011, https://perma.cc/SLR2-7LQY (keeping secretly recorded self-talk from evidence).

defendant's diary-like notes into evidence, on the reasoning that by setting his thoughts down in writing, he had released them from the inner sphere he controlled and exposed them to the risk of access. The four judges who would have found a violation in the writings being admitted as evidence reasoned that the notes were exclusively of a highly personal character, like a soliloquy, and did not lose protection by being entrusted to paper.[58] German law thus holds both sides of this Article's boundary simultaneously: Audience-less spoken thinking is absolutely protected, but the same thinking, typed, is balanceable because the line of absolute protection is drawn at the act of recording rather than at an audience. A chatbot exchange is often typed self-talk, sitting exactly on that fault line, and the German jurisprudence demonstrates two things at once: that a constitutional order can extend absolute protection to externalized thinking when it remains audience-less and that the recording-based line drawn by the *Tagebuch* plurality produces the arbitrary result that the constitution protects a thought spoken to a wall but abandons the same thought typed to a machine—the very premise this Article contests.

The same foundation exists in European Union law. CFREU Article 1 makes human dignity inviolable, and one settled strand of that guarantee, which Sue Anne Teo places at the center of her account on dignity and AI, holds that a person may not be treated as merely a means to someone else's ends.[59]

The production order described in the Introduction is that treatment. The conversations of 20 million people became evidence in a dispute they knew nothing about and had no part in, and their thinking was turned into material for someone else's case. That is a dignity problem, not just a privacy one, and it is why the standard I propose in Part V asks for more than the balancing that produced the order has delivered so far.

G. Neural Data: Neurorights Principles Build an Intermediate Tier

The neurorights movement has produced the most recent recognition that a specific category of data warrants protection because of its proximity to thought.

---

[58] BVerfG, 2 BvR 1062/87, 80 BVerfGE 367, 376, Sept. 14, 1989, https://perma.cc/5N76-T267 (*Tagebuch*). Because the judges split evenly, no violation of fundamental rights could be established.

[59] Sue Anne Teo, *Human Dignity and AI: Mapping the Contours and Utility of Human Dignity in Addressing Challenges Presented by AI*, 15 LAW, INNOVATION & TECH. 241 (2023).

Marcello Ienca and Roberto Andorno's proposal of rights to cognitive liberty, mental privacy, mental integrity, and psychological continuity has framed this debate.[60] Chile constitutionalized the protection of brain activity and the information derived from it in 2021,[61] and its Supreme Court applied the provision in 2023, ordering a consumer neurotechnology company to delete brainwave data collected through a commercial headset.[62] In the U.S., Colorado and California amended their consumer privacy statutes in 2024 to classify neural data as sensitive data that requires heightened treatment.[63] At the international level, the United Nations Human Rights Council requested guidelines on neurotechnology and human rights in April 2025,[64] and the United Nations Educational, Scientific and Cultural Organization (UNESCO) adopted its Recommendation on the Ethics of Neurotechnology in November 2025, articulating protections for mental privacy and cognitive liberty.[65] These neurorights principles are specifically interesting for the method they put forward: Confronted with data too revealing for ordinary treatment and plainly not inner thought, legislators built an intermediate category of heightened, qualified protection. That is the forum interstitiale under another name, constructed for the narrow case of data read from the brain, while the far larger stream of data typed out of the brain flows unprotected beside it. As an analogy, the neurorights instruments are thus interesting from a methodological point of view, but fail in terms of the object they cover: The protection they provide attaches to a data type defined by its physiological source, not to the function of the thinking act, and the related instruments never mention thought or conscience at all.

---

[60] Marcello Ienca & Roberto Andorno, *Towards New Human Rights in the Age of Neuroscience and Neurotechnology*, 13 LIFE SCIS. SOC'Y & POL'Y 5 (2017).

[61] Law No. 21383, Oct. 14, 2021, DIARIO OFICIAL [D.O.] (Chile), https://perma.cc/RZG6-GDFH (amending Article 19 of the Political Constitution of the Republic of Chile to protect brain activity and the information derived from it).

[62] Corte Suprema de Justicia [C.S.J.] [Supreme Court], Aug. 9, 2023, "Girardi c. Emotiv Inc.," Rol de la causa: 105065-2023, constitucional (Chile).

[63] H.B. 24-1058 (Colo. 2024) (amending the Colorado Privacy Act); S.B. 1223, ch. 887 (Cal. 2024) (amending the California Consumer Privacy Act to classify neural data as sensitive personal information).

[64] Human Rights Council Res. 58/6, Neurotechnology and Human Rights, U.N. Doc. A/HRC/RES/58/6 (Apr. 2, 2025) (adopted without a vote).

[65] UNESCO, Recommendation on the Ethics of Neurotechnology (Nov. 11, 2025), https://perma.cc/4G4K-VKZD.

H. The Convergence of Failures: Every Analogy Fails for a Similar Reason

Across privileges, query cases, reader protection, aggregation doctrine, the European law on correspondence and data protection, protections for drafts and private writings, and neurorights legislation, common characteristics can be identified.

First, a shared structure emerges. The protected act under freedom of thought is external, so the absolute protection of the forum internum is unavailable to that conduct. Ordinary expression doctrine is available, but it does not deliver what these acts need, and it falls short in different ways within each of the systems examined here. In the U.S., the First Amendment reaches private speech as well as public speech,[66] and the cases on protection for reading discussed above rest on that fact, but the Constitution supplies no general rule against compelled production of the records that speech leaves behind. The third-party doctrine and ordinary discovery relevance occupy that space instead.[67] Under ECHR Article 10 and CFREU Article 11, protection exists but is qualified, and the balancing is ordinary. In both the U.S. and the EU systems, the courts that did protect these acts reached past the ordinary expression standard for something stronger. The standard they reached for is neither absolutism nor ordinary balancing but a heightened, structured standard: scrupulous exactitude, demonstrated need, narrow standardized exceptions, partitioning, and human review. Table 1 sets out each analogy, the element it contributes, and the point at which it fails for the purposes of the interstitial zone.

---

[66] Speech on matters of purely private concern receives reduced protection in defamation and emotional distress cases. *See* Dun & Bradstreet, Inc. v. Greenmoss Builders, Inc., 472 U.S. 749, 758–61 (1985) (plurality opinion); Snyder v. Phelps, 562 U.S. 443, 451–52 (2011) (applying the public concern test and holding the speech at issue protected). Those cases calibrate tort liability for speech that has reached an audience. They do not condition First Amendment coverage on the existence of one, and neither addresses compelled production of the record that a thinking act leaves behind.

[67] On the third-party doctrine and the aggregation cases that have begun to limit it, see *supra* Part III.D.

Table 1: Each Candidate Analogy: What It Contributes and Where It Fails

Doctrine or principle	Protected act	Element contributed	Where the analogy fails
Seymour; *Tattered Cover* (*contra Kurtz*)	Search queries; purchase records	Compelling need; least intrusive means; showing made per person	Voluntary disclosure attack; one-shot act only
Jaffee and the dialogue privileges	Confiding to a professional	Predictability ex ante; narrow standardized exceptions	Requires a professional relationship; evidentiary rule only
Stanley; *Lamont*; reader privacy statutes	Reading and receiving ideas	Protection of the inputs to thought	Covers receipt, not the user's own contributions
Carpenter; *Chatrie*	Aggregated location records	Sensitivity to mosaics	Protects the pile, not the single revealing record
Digital Rights Ireland; *Tele2*; *La Quadrature du Net II*	Communications metadata	Partitioning, human review, strict necessity	Proportionality admits access on ordinary interests
ECHR Article 8; GDPR	Correspondence; special category personal data	Channel confidentiality; reuse and retention constraints	Correspondence presumes an addressee; data protection forbids little
Boyd lineage; work product; *Tagebuch* line	Private writings and drafts	Provisionality; the institutional parity argument	Content protection collapsed; line drawn at recording, not audience
Chile; Colorado; California; UNESCO	Neural data	The statutory intermediate tier as method	Attaches to brain data, but not to typed thought

The second finding that spans this table is the convergence of the analogies' failures. Each analogy breaks somewhere different on its surface: on the relationship (privilege), the disclosure (the search query post-*Kurtz*), the recording (the

Tagebuch line), the aggregate (the mosaic cases), the addressee (correspondence), the co-produced text (reading), the litigation shadow (work product), or the data type (neurorights). Yet each of these is the same failure seen from a different angle. In each of these doctrines, protection attaches to some feature that is incidental to thinking, and never to the act's function as thinking. No analogy fully works because none was built for the thinking process as such; that is why the interstitial zone must be constructed rather than borrowed, and why its boundaries must be the functional features of Part II rather than any doctrine's proxy. The fragments have never been assembled, in my view, because each of these doctrines developed separately and in response to a different technology. Human-AI interaction is what now makes the assembly necessary, since a single conversation with a general-purpose AI model collapses into one act what search queries, reading, confidences, and drafts once did separately.

IV. THE U.S. COURTROOM AND THE EUROPEAN UNION RULEBOOK

A. 20 Million Conversations Under Ordinary Balancing

The U.S. copyright litigation against OpenAI produced the first mass judicial treatment of conversational AI records. The plaintiffs sought a sample of 120 million ChatGPT conversation logs, and OpenAI countered with 20 million. Magistrate Judge Ona Wang ordered the full log sample to be produced, de-identified and subject to a protective order. On January 5, 2026, Judge Stein approved, holding that the magistrate judge had adequately balanced ChatGPT users' privacy interests against the relevance of the documents in light of the privacy protections already in place.[68] Nothing in that sequence is legally irregular. No rule was stretched and no judge erred: The logs were produced to the court by the ordinary operation of relevance and proportionality. Under existing categories, the conversations of 20 million people with a system many of them use as a confidant are deemed relevant commercial records, protected by de-identification and a protective order, and available on the balancing standard that governs any corporate document. No party to the dispute had a doctrinal name for the intuition, voiced loudly by activists and

[68] *In re* OpenAI, Inc., Copyright Infringement Litig., No. 25-md-03143 (S.D.N.Y. Jan. 5, 2026). Judge Stein wrote that "Judge Wang's rulings were neither clearly erroneous nor contrary to law. She adequately balanced ChatGPT users' privacy interests against the relevance of the documents in light of the privacy protections already in place." *See* Amanda Robert, *ChatGPT Creator Must Turn Over 20M Chat Logs in Copyright Litigation, Federal Judge Says*, A.B.A. J. (Jan. 8, 2026, at 11:13 CT), https://perma.cc/3BAC-EPM6.

other commentators outside the courtroom, that something in this record class differs in kind from a sales database. Nor do the safeguards reach that intuition. De-identification limits who can be attached to a conversation, not whether the conversation is read, and a person hesitates before asking a question if they think there is some likelihood of that question being connected to them.

B. Heppner *Against* Warner*: The Privilege Debate Asks the Wrong Question*

The individual-scale question arrived one month later in two U.S. court decisions issued in the same week. In *United States v. Heppner*, Judge Rakoff of the Southern District of New York held, on a question of first impression, that 31 documents a criminal defendant had generated by querying the consumer version of Claude were protected neither by attorney-client privilege nor by the work product doctrine.[69] Claude is not an attorney and recognized privileges require a trusting human relationship with a professional who owes enforceable duties. Moreover, the platform's consumer terms, which reserved rights to review human-chatbot interactions and disclose them, defeated any reasonable expectation of confidentiality. Finally, the court decided that materials prepared neither by nor at the direction of counsel are not work product, although the court left open whether direction by counsel might change the outcome of that question.[70] In *Warner v. Gilbarco*, a case decided the same day as the *Heppner* bench ruling's issuance, a magistrate judge in Michigan reached the opposite result for a self-represented civil plaintiff, denying a motion to compel her ChatGPT queries and responses. The judge held that generative AI programs are tools, not persons; that the materials reflected the thought processes of a litigant representing herself; and that work product protection is waived only by disclosure to an adversary or in a manner likely to reach one.[71]

The split has been analyzed principally by U.S. scholars and practitioners as a disagreement about waiver.[72] While that is an accurate depiction, for the purposes

---

[69] United States v. Heppner, No. 25-cr-00503-JSR, 2026 WL 436479 (S.D.N.Y. Feb. 17, 2026).

[70] *Id.* For further analysis, see Justin Curl & Mihir Kshirsagar, *AI and Privilege After United States v. Heppner*, LAWFARE (Mar. 30, 2026, at 13:00 PT), https://perma.cc/799X-MW4S.

[71] Warner v. Gilbarco, Inc., No. 2:24-cv-12333, 2026 WL 373043 (E.D. Mich. Feb. 10, 2026).

[72] *See, e.g.*, Sven Volkmer, Aalok Sharma & Hope Anderson, *Attorney-Client Privilege and Work Product in the Age of Generative AI*, WHITE & CASE LLP (Apr. 23, 2026), https://perma.cc/N7KB-HK68; Robert A. Schwinger, *In Initial Rulings, US Federal Courts Split on the Intersection of Litigant AI Use and Privilege Protections*, INT'L BAR ASS'N (Apr. 14, 2026), https://perma.cc/S7Z3-4TEG.

of this Article, the divide is relevant as a disagreement about classification. *Heppner* classifies the human-AI exchange relationally, asking what categorical bucket the AI falls into and concluding that, because it is not a lawyer, it is a third party and therefore a confidentiality-destroying recipient. *Warner* classifies the same relationship functionally, asking what the exchange does, and deeming the fact that it externalizes the litigant's own analysis an indication that it is an extension of her thinking, triggering the protection of her mental impressions. The relational frame seems to condemn every consumer-AI interaction in advance, because no mainstream AI model provider's terms establish professional duties. The functional frame asks a different question. It looks at what the person was doing, not at who or what was on the other side, and treats a person's exploratory dialogue as her own cognition in external form, whatever the tool and whoever is operating it. That is why the functional frame reaches the ordinary case and the relational frame cannot. Table 2 sets the two decisions side by side.

Table 2: Two Frames for the Same Act: *Heppner* Against *Warner*

	***Heppner* (S.D.N.Y.)**	***Warner* (E.D. Mich.)**
Posture	Criminal, represented defendant	Civil, self-represented plaintiff
Frame	Relational: what is the AI	Functional: what does the exchange do
Reasoning	Claude is not an attorney; consumer terms defeat confidentiality	Generative AI programs are tools, not persons; the materials reflect the litigant's own thought processes
Outcome	Production of all documents ordered	Motion to compel denied; work product protected
If generalized	Condemns consumer-AI interaction wholesale	Protects the human-AI exchange without personifying the machine

The European Commission has taken a similar functional position elsewhere. Asked in Parliament about the European Media Freedom Act, it answered that human rights are inherent to human beings, that it is always individuals who may avail themselves of expression rights, and that automatically generated content does not in itself enjoy such protection. This answer is not legally binding, but it shows the Commission drawing the line where the functional frame draws it: at the person,

not at the system she used.[73] If the machine's output carries no rights of its own, whatever protection the human-AI exchange attracts must be located in the human participant. The functional frame's premise has, in other words, already been stated by an EU regulator; only its protective consequence remains undrawn. The scholarly debate has divided along the same line. Sam Altman's call for an AI privilege, on the theory that talking to an AI should be like talking to a lawyer or a doctor, framed the protective case relationally and drew the predictable relational answer.[74] Ira Robbins identifies four criteria shared by the recognized privileges (a trusting human relationship, an enforceable duty of confidentiality, communications undertaken for the protected purposes, and a public interest that outweighs the evidential cost of granting protection) and concludes that AI interaction fails all four because the law protects relationships and not functionalities.[75] Within the relational frame, Robbins is correct, and commentary criticizing *Heppner* for overbreadth has struggled to escape the frame's gravity.[76]

There is a further objection, independent of Robbins's conclusion that no privilege is available to human-AI interaction. Privilege is an evidentiary rule. It operates at the moment a court compels testimony or production, and nowhere else. Even if an AI privilege were recognized and developed by the courts, it would have no bearing on the AI model provider's retention of the chatbot-user conversations, their use for model training, their disclosure under the provider's own terms, their acquisition outside a judicial process, or their pre-emptive filtering. The protection would begin and end at the courtroom door, and it would say nothing at all about the second direction of the problem, where interference in the thinking loop comes not from a record being read but from an answer being withheld.

---

[73] Answer given by the European Commission to Written Question P-001160/2025 (2025), https://perma.cc/7KHK-VAG6. *See also* NAJA BENTZEN, EUR. PARLIAMENTARY RSCH. SERV., INFORMATION MANIPULATION IN THE AGE OF GENERATIVE ARTIFICIAL INTELLIGENCE (2025), https://perma.cc/WVK3-EQH9 (reporting the European Commission's clarification, in answer to a parliamentary question on the European Media Freedom Act, that automatically generated and published content does not in itself enjoy the protection of free expression rights).

[74] Altman, *supra* note 5.

[75] Ira P. Robbins, *Against an AI Privilege*, HARV. J.L. & TECH. DIG. (Nov. 7, 2025), https://perma.cc/B7J2-YCZC.

[76] *See* Elizabeth X. Guo, United States v. Heppner, HARV. L. REV. BLOG (Mar. 23, 2026), https://perma.cc/A6QR-SGCE; Curl & Kshirsagar, *supra* note 70.

Moreover, Altman's support for the privilege proposal should not be interpreted as incidental. A privilege championed by the provider hosting the largest conversational corpus would function commercially as a confidentiality warranty attached to a product, encouraging users to confide in the system while the operative terms of service continue to reserve for the AI model provider the very access that *Heppner* found fatal to confidentiality. Whatever the sincerity of Altman's call for a privilege, the resulting structure would protect the relationship between the user and the product against third parties, and would leave the relationship between the user and the AI model provider untouched. But a protective doctrine built for freedom of thought must run in both directions: It must bind both the provider and the third party.

The interest worth protecting in a person's exchange with a machine is not the machine's professional status: It is the person's freedom of thought. From that perspective, what the machine is does not matter in terms of the granting of protection. Nobody asks whether a notebook owes its owner a duty of confidence before protecting a diary.

C. Thought Moderation: Filtering as Interference with Formation

Compelled disclosure is one direction of interference in the thinking loop. The other is pre-emptive restriction of what the mind may explore at all.

The European Union's AI Act prohibits, among other practices, the placing on the market or use of AI systems deploying subliminal, purposefully manipulative, or deceptive techniques that materially distort user behavior and cause significant harm.[77] Those prohibitions respond to the thought manipulation threat documented by human rights literature: Systems engineered to alter the thoughts of their users interfere with the forum internum from the inside.[78] The line that the literature attempts to draw is the right one: Influence becomes interference when it bypasses the user's capacity to reflect and resist rather than engaging it. This interfer-

---

[77] Regulation 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonised Rules on Artificial Intelligence, art. 5, 2024 O.J. (L 2024/1689), https://perma.cc/58VQ-C5LN.

[78] *See, e.g.*, SUSIE ALEGRE, FREEDOM TO THINK: THE LONG STRUGGLE TO LIBERATE OUR MINDS (2022); NITA A. FARAHANY, THE BATTLE FOR YOUR BRAIN: DEFENDING THE RIGHT TO THINK FREELY IN THE AGE OF NEUROTECHNOLOGY (2023).

ence is not just theoretical: It has been measured. In controlled experiments, participants who drafted text with an opinionated language model shifted their expressed views toward the model's position, which is what the loop model in Part II predicts. The effect was pervasive in the model's ordinary mechanics of writing assistance. No technique was used that a prohibition could name.[79] That is the difficulty with the prohibition model, which presupposes that the problem stems from a technique that can be pinpointed and forbidden. Here, the influence was a feature of the system doing what it was built to do—namely, to produce fluent continuations that carry a position. Banning models from holding positions would mean banning the assistance itself, and a rule that reached every system capable of shifting a view would reach every system capable of helping a person think. The line the AI Act draws against manipulative behavior by an AI system is the right one. What fails is its apparent assumption that interference always presents as a nameable technique.

This second direction of interference is less discussed. Compliance-driven output filtering, which I examined when the second draft of the General-Purpose AI Code of Practice appeared,[80] does not operate through any technique at all. Where regulatory instruments push AI model providers toward filtering model outputs for legal risk, that filtering operates on output that is generally produced one-to-one for a single user, and so on thought generation that is private by default. As I argued at the time, this inverts the logic of upload filters that have been used to moderate the upload of user-generated content on online platforms in a more dangerous direction, because it polices private thought experiments rather than public distribution of speech. Such types of output filtering are also contestable, and the Code's own history proves it: Its third draft deleted the model-level output similarity measures my analysis had targeted, and the final Code, published in July 2025, retained only a softened commitment on mitigating memorization, which refers to the tendency of a model to reproduce passages of its training data verbatim in its outputs.[81]

---

[79] Maurice Jakesch et al., *Co-Writing with Opinionated Language Models Affects Users' Views*, *in* 2023 PROCS. CHI CONF. ON HUM. FACTORS COMPUTING SYS. 111 (2023).

[80] De Cock, *supra* note 9.

[81] EUR. COMM'N, GENERAL-PURPOSE AI CODE OF PRACTICE (July 31, 2026), https://perma.cc/Y75A-ZZZ3. The second draft measure and its analysis are at De Cock, *supra* note 9.

I chose the term thought moderation deliberately, as an extension of the well-established concept of content moderation. Content moderation is the apparatus an entity puts in place to decide what may reach an audience, for its own purposes or under the pressure of policymakers or society. Thought moderation is the same apparatus, pointed at content with no audience. The techniques transfer wholesale to the AI model context: the implementation of filters on what people can receive as answers to their prompts and refusals to answer user requests that touch moderated boundaries, due to legal constraints or for the sake of mitigating liability or reputational risks. The justification does not transfer, however, because every rationale for moderating content (harm to audiences, distribution at scale, effects on third parties) presupposes the public availability that private generation lacks. The refusal of an AI model to answer a prompt with requested information or a filter on the information it is able to give in response to a query is therefore not a smaller version of content moderation: It is a different act with a different object, performed on the formation of thought rather than on its dissemination, and it inherits none of the legitimacy that moderation doctrine has built for the public sphere.

There is a second difference between the moderation of content and that of thought, and it matters more for remedies than the first. Content moderation is visible: A post comes down, an account is suspended, and the user is told, which is what makes appeal, oversight, and empirical study of such policing possible at all. Thought moderation, by contrast, often leaves no trace. Outright refusal of a chatbot to answer a prompt is the primary outlier to this principle and has the merit of being visible to the user who is denied an answer. In other cases, the AI system hedges, steers away, declines to develop a line of reasoning, or returns a thinner answer than the same question differently phrased would have drawn. The person receives an answer and never learns that she has been subjected to a restriction. She cannot contest what she cannot see, a regulator cannot count it, and a researcher cannot measure it. A person whose exploratory prompt is refused has not been prevented from speaking. She has been prevented from thinking with the instrument, and the classical binary between internal and external thought supplies no vocabulary for the harm. The forum interstitiale supplies it: Interference with interstitial acts requires justification at the heightened standard, whether the interfering hand belongs to a prosecutor seeking logs or to an AI model compliance system whose guardrails narrow user exploration in advance, visibly or otherwise.

The Future of Free Speech, which has tracked AI model refusals since 2024, found in its first study that the industry-leading chatbots declined more than 40 percent of prompts on lawful but contested political topics and that they would generate content supporting one side of an argument while declining the other.[82] Its more comprehensive follow-up found that blunt refusal rates have fallen since, but that models continue to steer around engaging in lawful discussion of reproductive rights, race, and armed conflict, and that the remaining refusals are justified by reference to acceptable use policies drafted in general terms rather than to comply with any stated rights standard.[83] Selectivity by an AI model in its answer, unlike outright refusal, leaves its user with no sign that anything was declined, which makes it the harder of the two problems even as the overall number of refusals improves. The Oversight Board, the independent body established by Meta to review its content-moderation decisions, measured that selectivity at scale in its first structured evaluation of large language models, published in July 2026. The report found the models were more than twice as likely to refuse to produce material critical of governments that restrict speech than of governments that do not, refusing 34 percent of requests concerning restrictive jurisdictions, but only 14 percent for permissive ones.[84] The incentive behind that caution is easy to trace. In May 2026, a German court held Google liable for the content of its AI-generated summaries.[85] Once an AI model provider can be sued for what its system produces, refusing is cheaper than answering. And the people writing the rules that produce that incentive are not thinking about thought at all. The European Parliament's own research service briefed Members on generative AI and information manipulation, covering

---

[82] Jacob Mchangama & Jordi Calvet-Bademunt, *Freedom of Expression in Generative AI: A Snapshot of Content Policies*, FUTURE OF FREE SPEECH (Mar. 1, 2024), https://perma.cc/R6V4-Z2CY.

[83] Jordi Calvet-Bademunt, Jacob Mchangama & Isabelle Anzabi, *That Violates My Policies: AI Laws, Chatbots, and the Future of Expression*, FUTURE OF FREE SPEECH (Oct. 9, 2025), https://perma.cc/RJL7-2Z35. The study also ranks six jurisdictions by how their AI laws treat lawful but controversial expression, which is the regulatory direction of the threat model assessed comparatively.

[84] OVERSIGHT BD., ARE LLMS STIFLING POLITICAL SPEECH? AN ASSESSMENT OF HOW AI MODELS PROTECT FREE EXPRESSION (2026), https://perma.cc/YY3H-GCXM.

[85] Katharina Wagner & Jenny Gesley, *Germany: Court Holds Google Liable for Incorrect AI Overviews*, LIBR. OF CONG.: GLOB. LEGAL MONITOR (July 17, 2026), https://perma.cc/7LZ3-TQPP.

the AI Act, the Digital Services Act, and the European Media Freedom Act. The briefing did not mention the right to freedom of thought once.[86]

V. The Forum Interstitiale: Definition, Boundaries, and Protection Standard

A. Definition and Scope: The Anteroom Between the Fora

The preceding Parts have supplied the materials: a theory of why some external acts are thought-like, an analysis showing that every candidate analogy protects part of the thinking process but fails to satisfy here because it does not protect the entire thinking process, and a body of case law demonstrating what happens when the intermediate category of thought has no name. Since no analogy can be borrowed whole, the category has to be built rather than adopted, out of the elements contributed by the analogies reviewed above but attached this time to the function they all missed. The forum interstitiale is the zone of external acts that perform the function of thought formation without constituting communication to a public. An act belongs in the zone when it satisfies the functional profile developed in Part II: It is provisional rather than final; its intent is non-communicative or instrumentally dialogic; no human reader receives it as a finished communication, and any who read it are participants in its formation rather than its destination; and it serves epistemic exploration. A silent reading of generated text satisfies all four requirements fully and sits closest to the forum internum. A near-final draft shared with a co-author begins to fail the audience criterion and slides toward the forum externum, where it is protected as expression rather than as thought. Acts within the zone therefore do not all attract the same protection, and protection falls as human circulation widens. The taxonomy in Part VI sets out which acts sit closer to the forum internum and which are more adjacent to the forum externum. Table 3 labels the first group high interstitial and the second middle interstitial.

A room metaphor runs through my wider work on these questions. If the forum internum is the room inside your head and the forum externum is the public square, the forum interstitiale is the anteroom between them: the space where a thought is prepared for presentation by its owner before its public debut, and where it can be left hanging if its owner decides to return into internal solitude for further contemplation. Figure 1 shows the three fora, the two directions of pressure on the interstitial zone, and the no-demotion rule.

---

[86] BENTZEN, *supra* note 73.

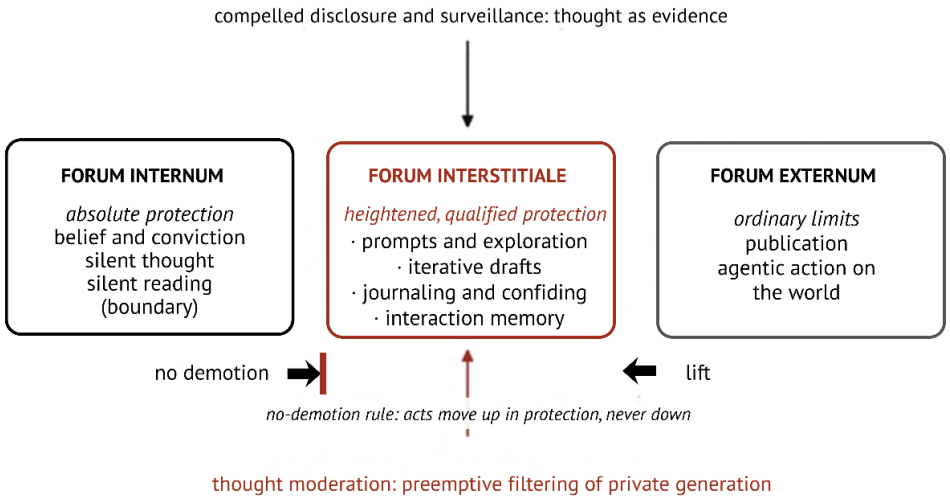


Figure 1: The Three Fora: The Interstitial Zone Under Pressure from Both Directions, and the No-Demotion Rule.

The definition of this liminal space deliberately attaches to acts, not technologies. The interstitial zone covers the search query and the paper diary as readily as the chatbot conversation, which is why the doctrines assembled in Part III could develop before conversational AI existed. What AI changes is scale and consolidation: A general-purpose conversational AI system hosts, in one place and one record, the full range of interstitial activity a person previously might have distributed across searches, notebooks, drafts, and confidants. The scale of the information mosaic created makes apparent the urgency of generating appropriate protection for the acts implicated.

B. The Protection Standard: Five Elements

No legal system currently recognizes the interstitial zone I propose, and nothing currently receives protection on the basis that it belongs to that zone. What follows is my account of what that protection should be. Acts within the zone would receive heightened but qualified protection, structured by the elements already tested by the fragment doctrines. The five elements below form the structure of the proposal itself. The applications attached to them, and the placements in Table 3, are worked illustrations: They show that the standard can be operated on facts that courts are already deciding, and they indicate where each element points. They are not a draft code. Specific thresholds, exceptions, and evidentiary consequences

would each need justification of a kind this Article does not attempt, save for one that carries enough weight to require it, and to which I return in Part VI.

First, predictability. In *Jaffee*, the U.S. Supreme Court reasoned that a privilege whose application is weighed case by case is worth little, because the person deciding whether to speak must make that decision before any court is involved. The same holds here. A protection that depends on later balancing at the point of access cannot produce the candor and exploration it exists to enable, since the thinking that may be entitled to protection happens long before anyone asks for the record of it.

Second, a heightened access threshold. Compelled disclosure of the records of an interstitial act would require individualized suspicion and a demonstrated need, applied with the exactitude *Seymour* demands, rather than ordinary relevance balancing. Access to such records as a class, whether across many people or across a long period of one person's activity, would be presumptively impermissible. The 20-million-log production order is the founding example, and it is the case most directly reached by the aggregation reasoning: A collection assembled from many people's records, or covering a long stretch of one person's, reveals more than any of its parts and should not be available on a showing that would not support access to any single conversation in it.

Third, structural safeguards on any permitted access. The CJEU case law on communications metadata supplies three: identity data partitioned from content data, no fully automated review, and a human oversight decision before access. The privilege model supplies a fourth requirement of a different kind: Exceptions must be defined in advance of any interference rather than left to case-by-case discretion, since an exception whose scope is settled after the fact leaves the thinker seeking protection no better off than case-by-case balancing would. Something of the order of imminent harm to life is a possible candidate, though *Jaffee* left the question open.

Fourth, protection against pre-emptive restriction. Interference in the thinking process need not wait for a record to exist. Where compliance incentives drive a system to filter or decline private exploratory outputs, the user's exploration is cut short rather than recorded, and that interference would require justification at the same heightened standard as compelled disclosure. What matters in this assessment is whether the restricted output is private and formative or public and distributive. That heightened justification requires the interference to be visible, and the

evidence in Part IV.C is that it generally cannot be. The grounds for refusal sit in acceptable use policies drafted in general terms, and no AI model provider discloses the training and alignment decisions that determine what a model will engage with. The element therefore carries a second limb. Where a user's exploration is filtered or refused, they are entitled to notice that a restriction operated and to an explanation of the grounds on which it operated.

Fifth, the no-demotion rule. The interstitial zone only works upwards. It takes acts that today sit in the forum externum, or that receive no protection at all, and raises them. It must never be used the other way. Anything the law currently places inside the forum internum stays there, absolutely protected. The obvious way to misuse a middle tier is to demote things into it, and nothing in this proposal licenses that.

These five elements are addressed mainly to courts and governments, and mainly to what they may not do. That is not enough. Sue Anne Teo makes the point that this is the smaller half of freedom of thought: Forbidding violations does less work than requiring that the surrounding environment be built in such a way that ensures the freedom can be used in the first place.[87] None of the five elements works if the conversation was never private to begin with. An AI model provider whose terms let it read, keep, and reuse what its users explore has already given away the confidentiality that everything above depends on, and no rule of evidence can put it back. So part of the work falls on AI model providers: limiting retention of user information to what provision of the model's service requires, storing user identity data separately from content, and creating at least some interaction modes in which the provider does not retain readable records. How far any of these can be required, and through which instruments, is a question this Article leaves open.

C. The Dilution Objection: A Qualified Tier Protects the Absolute Core

The objection that expanding freedom of thought dilutes its absolute core is the most serious in the field, and it is a valid objection when assessed against the proposals it has so far targeted. A right that absolutely protects everything functionally connected to cognition would collapse at first contact with a compelling countervailing interest. The forum interstitiale is not such a proposal. It expands nothing

---

[87] Sue Anne Teo, *How to Think About Freedom of Thought (and Opinion) in the Age of AI*, 53 COMPUT. L. & SEC. REV. 105969 (2024).

that is absolute and creates a qualified tier so the absolute tier need not stretch, relieving the pressure for it to do so. Where McCarthy-Jones and Walmsley extend conditional absolute protection into the forum externum and must then defend the absoluteness of the extension, the interstitial architecture accepts that external acts can be limited and sets the conditions for limiting them.[88] The dilution objection does not defeat this proposal. It explains why the middle tier must exist, and why it must be qualified.

VI. Mapping Individual AI Actions Across the Three Fora

Table 3 maps the principal actions an individual performs with AI systems across the three fora. Placements follow the boundaries of Part V and set out the supporting analogies and the direction in which the standard would point.

Table 3: Individual AI Actions Mapped Across the Three Fora

Action	Classification	Supporting analogy	Indicative consequence
Silent reading of AI output	Forum internum boundary; highest interstitial	*Stanley*; *Lamont*; reader privacy statutes	Near-absolute; state inquiry into what a person reads is subject to the most demanding justification
Information-seeking prompt	Interstitial, high	*Seymour*; intellectual privacy of the query	Access only on individualized suspicion, applied with scrupulous exactitude
Exploratory or hypothetical prompting	Interstitial, high	Extended thought; thinker-based theory; provisionality	Shielded as pre-committal exploration; content of hypotheses not sufficient on its own to prove belief or intent
Journaling, venting, therapy-like conversation	Interstitial, high	*Jaffee*; the confiding strand of intellectual privacy	Qualified privilege; disclosure only under narrow standardized exceptions such as imminent harm to life

[88] McCarthy-Jones & Walmsley, *supra* note 18.

Action	Classification	Supporting analogy	Indicative consequence
Iterative drafting and brainstorming	Interstitial, middle	Private writings; work product; *Warner*	Protected as mental impressions; tool use is not waiver; circulation for comment reduces protection without ending it
AI memory and persistent interaction profiles	Interstitial, middle, aggregation-sensitive	*Carpenter*; *Chatrie*; CJEU metadata line	Bulk access presumptively impermissible; partitioning and human review; access scales with what the aggregate reveals
Agentic delegation acting on the world	Forum externum	Manifestation and conduct doctrine	Ordinary regulation; effects on third parties justify oversight of the action, though not retrospective harvesting of the formative dialogue behind it
Publishing AI-assisted output	Forum externum	Freedom of expression	Ordinary expression rules and their qualified limits

Three placements need further explanation. The exploratory prompt sits where the standard bites hardest, because it is the act most exposed by the binary: Search history is already routinely offered as evidence of the searcher's intent, and a doctrine that reads the devil's advocate prompt as an admission converts the testing of ideas into self-incrimination, producing the chilling effect Penney measured in the wake of the Snowden disclosures. Agentic delegation marks the zone's outer edge and its most contested one. When a person instructs an AI agent to act in the world, the action joins the forum externum, but the dialogue that preceded the instruction does not follow it there, any more than signing a contract renders discoverable the diary in which the signing was contemplated. The AI memory and profiles debate is often linked to the aggregation point, because persistent profiles are where the mosaic reasoning applies most forcefully: A system that remembers a year of a person's exploration holds a longitudinal record of a mind, and access to that chronicle should be assessed against the whole rather than the fragment.

The exploratory prompt needs more than that level of protection. My claim is that it should not by itself be deemed to prove what a person believed or intended, and my reason for that position is based in ordinary evidence law rather than in anything new. Someone who is exploring an argument and asks a machine to argue one side of an argument has not necessarily adopted that side, so the step from asking to believing is a weak one. In contrast, a prompt read back in a courtroom has a persuasive effect on a judge or jury out of proportion to what it actually shows. That combination—a weak inferential connection and a strong prejudicial effect—is already a recognized ground of exclusion, and it applies here. The result is not immunity. A prompt can still corroborate intent that other evidence establishes; what it cannot do is stand alone.

VII. Objections and Answers

A. Why Not Simply Repair the Binary?

The foremost objection to the argument laid out here is that none of this is necessary. Courts confronting these acts, as Part IV shows, have reached for privacy and evidentiary doctrine rather than for speech doctrine. If the difficulty is therefore that they cannot tell an interference with privacy from an interference with thought, the answer might be a better test rather than a new category.

Sue Anne Teo gives the most developed version of that answer. She does not attack the binary; she works inside it. She asks when an interference stops being a privacy problem, which courts can balance, and becomes a freedom of thought problem, which they cannot, and she locates the dividing line at the point where a practice damages what people need in order to think for themselves. She adds that restraint is not enough: Whoever designs the environment has to build it so that people can still think in it.[89]

Where I differ is on why the line is so hard to draw. The binary sorts internal beliefs from their external manifestation, and a search query is sometimes neither. It is not necessarily a belief, and it is not always the manifestation of one. No refinement of the test will place it, because the scheme has no place for it to go. Teo has to do a great deal of work to locate a spot for the dividing line to rest, and I take the difficulty as evidence that such a line cannot resolve all of the problems inherent to the binary's structure, rather than as a reason to keep looking for where that line might rest more easily.

---

[89] Teo, *supra* note 87.

B. Line-Drawing: Borrowed Boundaries, Not New Ones

Any intermediate category invites the objection that its boundaries will be arbitrary. The first answer to that challenge is that these boundaries are borrowed rather than invented: the audience test from work product and privilege law, the provisionality test from the law of drafts, the aggregation trigger from *Carpenter* and the CJEU metadata cases, the exactitude standard from *Seymour*. Each has already been applied and tested in litigation. The second answer is that the alternative is not a sharper boundary but none at all. The same kind of ChatGPT exchange was protected work product in Michigan and discoverable in Manhattan in the same week, which is what happens when courts have no category to reach for. A named category with stated criteria gives them a test they can apply consistently. An unnamed gap guarantees they will keep reaching different answers.

C. The Relational Objection: Privileges Protect Functions, Not Relationships

Robbins is right about the privileges as they stand.[90] They attach to relationships: lawyer and client, doctor and patient, priest and penitent. A chatbot is none of these. But that tells us how the law arrived where it is, not where it should go.

The consulting room is not protected because therapists deserve something. It is protected because patients will not speak freely if they think their words can be repeated outside that safe space, including in court. The professional relationship is how the law has guaranteed that those words stay private. It is the mechanism, not the reason.

People now speak equally freely to machines. The reason for protecting the conversation is intact, but the mechanism is missing. The answer to a missing mechanism is to build one, not to conclude that the conversation does not matter. That is what Part V asks of AI model providers: Keep less, separate identity from content, and offer at least some modes in which nothing readable is kept.

The *Heppner* case illustrates the point. The New York court relied heavily on terms of service that let the provider read and disclose what users wrote. Read that way, its decision is less a ruling about privilege than a verdict on how AI model providers handle data. AI model providers who want their users' exchanges to be protectable will most likely have to make them confidential first.

---

[90] Robbins, *supra* note 75.

D. Does Everything Fall Under Freedom of Thought? The Exclusion Criteria

If querying a chatbot is thought-like, why not every use of every tool? The answer is in the interstitial zone's exclusions. An act that has reached an audience is out. So is a position that has been settled, adopted, or communicated, and an action taken on the world. Between those exclusions and the core of the protected forum interstitiale, an act circulated for comment but not yet finished sits lower in the zone, closer to the forum externum. The zone captures the stage where the thought is formed and nothing else. The objection that acts undertaken with tools beyond chatbots should be included has real force against proposals that extend the forum internum's absolute protection to whatever is functionally integrated with thinking, because absolute protection tolerates no vague edge. It has much less force against a qualified tier. An act wrongly placed inside the zone does not hand the person who committed that act an inviolable shield; it means only that the state must offer more justification for any interference with the act than it otherwise would.

E. The Seriousness Threshold: Should Mundane Prompts Even Qualify?

The European Court of Human Rights requires that a belief attain "a certain level of cogency, seriousness, cohesion and importance" before it engages the protection of ECHR Article 9.[91] A critic may object that the throwaway prompt and the abandoned draft fall far below that threshold, and that a category built to hold them protects nothing the right to freedom of thought was ever concerned with. That right's threshold, however, was developed to identify protected beliefs for the purposes of manifestation claims, not to filter the acts through which beliefs are formed, and the transposition from one to the other is not self-evident. Teo observes that the distinction between the trivial and the consequential is precisely what inference dissolves: A system that draws sensitive conclusions from innocuous data makes illusory any sorting of mental content by apparent importance.[92] The same holds for the interstitial record, where the significance of a single query is often visible only in aggregate or in retrospect. The deeper answer is that this seriousness objection is an argument about protecting the absolute tier, where a demanding threshold for entrance is the price of the absolute protection an act receives after its entry. The forum interstitiale is qualified and can therefore afford to

---

[91] Campbell v. United Kingdom, 4 Eur. H.R. Rep. 293, ¶ 36 (1982).

[92] Teo, *supra* note 87.

admit the mundane. What attaches to an idle prompt is not inviolability but a requirement that the state justify itself before it may compel production.

F. The Downgrade Risk and the No-Demotion Rule

The most serious residual danger is the one addressed by the no-demotion rule: that a named middle tier becomes a destination to which courts demote acts previously held in the forum internum. The risk is real, and a rule cannot by itself stop a court that is determined to ignore it. What the rule can do is remove the route for demotion. The interstitial zone admits only external acts, so anything the law has already placed inside the forum internum does not satisfy the entry criteria for the intermediate zone and cannot be moved there without those criteria being abandoned. Silent belief and unexpressed conviction are therefore not candidates for its inclusion. Set against that risk of demotion is what is happening now. Without a middle tier, the acts that I contend belong in the forum interstitiale are not being held safely in the forum internum awaiting demotion: They are being processed as commercial records at the rate of 20 million conversations per order.

Conclusion: The Need to Fill the Gap

The forum internum and forum externum were adequate categories for a world in which thinking was invisible and only its expression left traces. That world has now evolved into something else entirely. Today, thinking leaves records, because the instruments supporting thought are external, hosted, and logged, and the law's binary silently converts the records of formation into the most exposed class of documents a person generates. The results are already visible. Conversations have been compelled by the tens of millions under ordinary balancing, criminal defendants have had their exploratory dialogues produced against them, and identical acts were protected in one district and stripped of protection in another within a single week. The same legal systems, meanwhile, extend heightened protection to search queries, reading records, confidential dialogue, location mosaics, and brainwaves, without noticing that the fragments they have built describe a single category. The U.S. debate about granting legal privilege to AI is the clearest symptom of the missing classification. Asked to protect the thinking of humans, the law reached for a doctrine about the duties of professionals, and the loudest voice proposing that analogy was one of the AI model providers hosting the conversations.

This Article has named the category, specified its boundaries, and assembled its doctrinal form from parts that the law has already built and tested. The forum interstitiale asks courts and legislators for less than the absolutist expansion debated

by the cognitive liberty literature, and for far more than the status quo delivers: a predictable, heightened, qualified shield over the external acts through which people now form their thoughts, bounded by a no-demotion rule that leaves the absolute core exactly where the drafters of the ICCPR placed it. What is at stake is not taxonomy but rather the practical conditions of free thought in the systems where thinking now happens; and those conditions will be settled in regulation, discovery orders, and terms of service, whether or not the settlement is given a name.

Litigators should argue for the functional frame *Warner* adopted, not the relational frame *Heppner* utilized to deny protection. Legislators should hold every new AI rule against one screening question: Does the measure push AI systems to manipulate thought, or push AI model providers to filter it, or neither? They should also write the interstitial standard into the instruments that currently treat conversational records as ordinary data. Crucially, AI model providers hold the piece without which none of this works. A rule of evidence cannot restore confidentiality that a terms of service page has already signed away, so the utility of this standard depends on providers keeping less information; separating identity from content; and being in a position to say truthfully, when a court asks, that there is nothing readable to hand over.

I would rather see this new category named and built on purpose than assembled by accident, one discovery order at a time, or not at all.