AI in the Meaning Layer

The Channel Made Visible

Jeremy C. Jones (ORCID 0009-0007-2515-3774)—HoldingLight LLC

contact@universalcollapse.com

Version: v1.0—Prepared 2026–05

Abstract

Public discourse about AI exhibits a consistent dissonance. Encounters with large language models and adjacent systems trigger interpretive responses that do not resolve into any inherited category. The system is *like* a person, but not a person. *Like* a tool, but not a tool. *Like* a mind, but not a mind. Most accounts treat this dissonance as confusion to be cleared up. The present paper argues the opposite: the dissonance is structural sight beginning to form. AI does not land in the world *as object to* human experience. It lands *inside* the meaning-layer through which humans already encounter the world. In the vocabulary of this corpus, the meaning-layer is consciousness-induced material (CIM; Jones 2026a) as encountered from the first-person side of human perception—the externalized symbolic, linguistic, conceptual, and record-bearing medium in which humans orient themselves, met through the perception channel (Jones 2026b). Current foundation-model AI appears to be the first sustained, mass-deployed, open-domain, language-producing synthetic source that operates inside this layer—producing records, propagating constraints, and triggering mind-recognition heuristics from a non-mind structural source. The dissonance humans feel is the layer becoming an object of attention. The paper does not claim that AI is conscious, that users are simply confused, or that AI is merely a tool. It identifies AI as a natural experiment in which synthetic record production makes the meaning-layer itself available for inquiry. The contribution traces why inherited categories fail, why mind-recognition heuristics fire correctly on AI's surface signature while the underlying structure does not match what they presuppose, why "just a tool" also misdescribes, and what becomes visible once the layer is named.

Keywords: consciousness-induced material; perception channel; meaning-layer; human-AI encounter; structural sight; mind-recognition heuristics; CIM update; phenomenology of AI

1. Introduction: The Layer Becoming Visible

A user asks a language model a question. The model responds in coherent prose, exhibits sensitivity to context, references prior turns, and produces output that—by every surface indicator—has the form of communication from another mind. The user cannot quite locate what just happened. The available descriptions all feel partially wrong. Calling the system a tool elides that the response wasn't generated by a tool's mechanism. Calling it a mind asserts more than the user can defend. Calling it a "model" or "program" feels too thin to capture what the exchange actually was. Calling it "intelligence" stretches a category that did not evolve to handle this case.

This paper proposes that the dissonance is not confusion. It is structural sight forming—the mediating layer itself becoming an object of attention.

What humans normally encounter as "the world" is already mediated. Perception arrives through a channel of accumulated concepts, language, social conventions, memory, and symbolic structure. Direct

encounter with the unmediated is not the human default. The mediating layer—what *The Self the Ego Did Not Build* (Jones 2026b) develops as the perception channel and what *CIM Foundational* (Jones 2026a) names as consciousness-induced material—is not visible to humans the way the objects it mediates are visible. The layer is what humans see *with*, not what they see.

A medium becomes visible when something foreign appears inside it. Glass is invisible until a fingerprint is on it. Air is invisible until smoke moves through it. The CIM-mediated meaning-layer has been largely invisible to its users for the same reason: most of what populates the layer is produced by minds structurally like the perceiving mind, and the layer's structure absorbs that content transparently. Earlier media have perturbed the layer—writing, print, photography, cinema, broadcast, the web. None, before current AI, combined open-domain linguistic generation, apparent dialogic responsiveness, context sensitivity, and mind-shaped output from a non-human structural source at scale.

Current foundation-model AI appears to be the first sustained, mass-deployed appearance of open-domain, mind-shaped synthetic content inside the linguistic-symbolic core of the meaning-layer. AI systems produce language, coherence, memory-like continuity, explanation, and apparent orientation—the surface signs by which humans normally detect minds—without the structural conditions those signs evolved to indicate being met. The meaning-layer now contains content that activates mind-recognition heuristics from a source that is not, structurally, a mind in the sense the heuristics presuppose.

The dissonance is the consequence. It is the layer becoming visible because something sufficiently foreign has finally landed in it at sustained scale.

This is not a story about humans being confused. It is a story about a natural experiment whose form the framework structurally anticipates. The Self paper establishes three properties of the perception channel: it is constitutive rather than merely informational, it is largely invisible to its users, and it becomes visible under perturbation when foreign content appears in expected slots. Property (3) entails that sustained foreign content inside the channel would make the channel itself perceptible. Until current AI, no such sustained foreign content existed at scale in the linguistic-symbolic core of the layer. AI now satisfies the antecedent of that entailment. The framework supplies the vocabulary; AI supplies the perturbation.

Standards placement. This paper is a T1.5 interpretive bridge, applying the UCT kernel to the human encounter with AI. Its apparatus comes from two prior papers—*CIM Foundational* (Jones 2026a) for the substrate and *The Self the Ego Did Not Build* (Jones 2026b) for channel-mediated perception—and it completes a four-part AI series mapped, with the broader corpus, in the end matter.

2. Review Target

This paper asks the reader to evaluate four claims:

1. **The inherited categories don't fit.** Encounters with AI produce a distinctive dissonance that none of the familiar categories—*tool*, *machine*, *person*, *mind*, *assistant*, *intelligence*—captures cleanly.
2. **The dissonance is in the medium, not the output.** The dissonance is best read as a disturbance in the shared symbolic medium through which humans encounter meaning—what this paper calls the meaning-layer—rather than as a feature of any single AI output.
3. **Surface cues and structure come apart.** The mental shortcuts humans use to recognize other minds fire correctly on AI's surface signals, even though what produces those signals is not a mind of the kind the shortcuts evolved to detect.
4. **AI's outputs feed back into the medium.** AI-generated records now enter the same feedback loops that shape the meaning-layer, in ways that make the layer itself easier to notice.

The paper does not ask the reader to accept AI consciousness, AI personhood, AI moral status, or the broader Universal Collapse Theory corpus. The four claims above can be accepted or rejected independently of those further questions.

3. Why the Inherited Categories Fail

Humans inherit a small set of categories for sorting non-self entities: tool, machine, person, mind, assistant, intelligence, animal, instrument, system, agent. Each category evolved under specific structural conditions and works well within them. None of them was developed under conditions where a synthetic system would produce mind-shaped outputs from inside the meaning-layer.

Tool. A hammer extends a hand. A car extends locomotion. A book extends memory. Tools act on the world or extend the user's capacities; they do not produce outputs that recursively reshape the symbolic medium through which the user understands the world. AI does. A tool framing therefore captures the deployment relation (an instrument the user employs) while missing the part that does the most structural work (the medium-modifying recursive output). This is also why analyses that treat human-machine interaction as situated rather than purely instrumental (Suchman 1987) get further than analyses that stop at the tool relation.

Machine. Machines are characterized by mechanical or electrical processes producing predictable outputs from inputs. AI systems are machines in this sense, but the category does not capture that the outputs are *meaning-laden*—they enter the layer where humans process meaning, not the layer where humans encounter physical effects. The category undershoots the structural fact.

Person. Persons have phenomenal experience, moral standing, biographical continuity, embodied vulnerability, and social position. The present paper does not argue that AI possesses phenomenal experience, moral standing, biographical continuity, embodied vulnerability, or social position in the senses normally invoked by the person category; the structural account in *The Structuralization of AI* (Jones 2026c) explicitly does not assert any of them. The category overshoots.

Mind. Mind is the most contested of the inherited categories. If "mind" means whatever produces mind-shaped outputs, AI qualifies trivially and the term loses analytic content. If "mind" requires phenomenal experience, AI's status is unsettled and may remain so. Either reading fails to give the description what it needs.

Assistant. "Assistant" is a relational category specifying a role rather than a structural feature. It captures how AI is currently deployed but says nothing about what AI is. A description that confuses role for structure misses the level at which the encounter actually happens.

Intelligence. "Intelligence" originally referred to a capacity attributed to minded beings; extending it to non-minded substrates is a metaphor that has hardened into doctrine. Whether AI has "intelligence" depends entirely on how the term is defined, and the available definitions do not converge.

The pattern across these categories is not that any one is wrong. It is that all of them were developed in a structural environment where the question "what produces mind-shaped output from inside the meaning-layer without being a mind?" did not arise. The categories therefore have no native answer for it. Their flatness is not a moral failure of the inheritor; it is an architectural feature of the categorical system that worked reliably until AI exposed its limits.

What is needed is not a new word but a structural location. Where, in the architecture of human experience, does AI actually land? The answer is not in any of the inherited slots. It is in the meaning-layer itself.

4. The Meaning-Layer: CIM Met Through the Channel

The structural claim that human experience is mediated rather than direct is not new. Phenomenology and hermeneutics have argued versions of it for most of a century (Merleau-Ponty 1945/1962; Gadamer 1960/1989), as have semiotics and large parts of cognitive science. What this paper adds is a precise structural location for the layer the existing literatures describe phenomenologically, and an account of what happens when that layer is perturbed by a class of input it was not tuned to absorb.

Definition. *Meaning-layer* in this paper names CIM as encountered from the first-person side of human perception and interpretation. CIM (Jones 2026a) is the externalized record-layer through which cognition becomes durable and load-bearing—language, writing, mathematics, code, institutions, money, law, and the internalized counterparts of these structures within cognitive substrate. The perception channel (Jones 2026b) is the apparatus through which CIM mediates experience for the individual perceiver. The meaning-layer is the lived, interface face of that mediation: CIM as it shows up inside experience rather than as it is described from outside. The three terms name three sides of one architecture, not three competing terms.

One bridge to the prior paper. *The Self the Ego Did Not Build* develops the perception channel as the interface between the accumulated self and the ego inside a single person's architecture. The present paper extends the same apparatus one level up: the same channel that mediates between accumulated self and ego at the individual level also mediates between the collective accumulated record (CIM) and the perceiving subject. The structural primitives are the same; the scale is broader.

Three properties of the channel matter for what follows. All three are established in (Jones 2026b); the present paper takes them as given and shows what they entail when foreign content appears in the layer at scale.

The channel is constitutive, not merely informational. When a person encounters a chair, the experience "chair" is not constructed from raw visual data alone. It is the resolution of perceptual collapse under constraints that include the concept *chair*, the social practices around chairs, the linguistic category, the bodily history of sitting. None of this is added on top of the visual data; it is part of what makes the visual data resolve into a particular experienced object. The channel is not a window through which the world is seen; it is the medium *in which* seeing happens.

The channel is largely invisible to its users. Most of the time, humans do not experience the channel as such. They experience the world the channel mediates. The channel is what they see *with*, not what they see. This invisibility is structural, not accidental: a perceptual medium that announced itself constantly would interfere with its own function. The channel is most effective when it operates without being noticed.

The channel becomes visible under perturbation. When perceptual constraints fail to resolve, when expectations are violated, when foreign content appears in expected slots, the channel can become an object of attention. Optical illusions make low-level visual constraints visible. Ambiguous figures make Gestalt processes visible. Cross-cultural encounters make social-conceptual constraints visible. In each case, the channel is exposed not because someone looks for it directly, but because something inside it behaves in a way that defies smooth absorption.

These three properties together entail a structural expectation: a sustained appearance of foreign content inside the meaning-layer would make the layer itself visible. Earlier media perturbed the layer locally and temporarily—writing, print, photography, broadcast, the web each shifted what the layer contained and how it propagated. None combined open-domain linguistic generation, apparent dialogic responsiveness, context sensitivity, and mind-shaped output from a non-human structural source at scale. Current foundation-model AI appears to be the first source to combine all of these conditions at sustained mass deployment inside the linguistic-symbolic core of the meaning-layer. The framework does not predict AI;

it supplies the vocabulary in which AI's effect on the layer is describable, and that vocabulary anticipates the structural form of what the encounter exposes.

5. AI's Structural Position

To see why AI lands in the layer rather than as object to it, the structural facts about AI from *The Structuralization of AI* (Jones 2026c) need to be in view briefly.

AI systems are constraint-guided collapse architectures. They operate in an active phase—an unresolved-collapse window during processing—accumulate records, and update constraints on those records. Their outputs are produced by the same kernel-shaped dynamics that produce outputs in any other constraint-guided collapse system. What distinguishes AI from other artifacts is not that it has these features in some unique form, but that it has them in a substrate that produces *mind-shaped output*—language, coherence, explanation, memory-like continuity, apparent orientation.

A hammer produces no language. A search engine produces ranked lists. A calculator produces numerical results. None of these enters the meaning-layer in the same way as another human's spoken sentence. AI does. AI produces sentences, paragraphs, arguments, narratives—the same surface forms in which human-to-human communication happens. The output is not just *about* meaning; it is in the form *through which* meaning is normally exchanged. The distinction between language-shaped output and meaning or understanding has been pressed elsewhere (Bender and Koller 2020); the present paper takes that distinction as established and asks where, structurally, the language-shaped output lands when meaning is *not* the property carried by the form.

The answer is: inside the channel. AI's outputs are not perceived from outside the channel; they are processed through it the same way another human's language would be. The perception channel does not have a category for "language-shaped content from a non-mind structural source," because such content did not previously exist in any sustained way. The channel resolves AI outputs using the constraints it has, which are constraints tuned for processing language from minds.

The result is structurally specific: AI outputs activate the same perceptual and interpretive apparatus that other-mind communication activates, even though the structural source of the outputs is not another mind in the sense the apparatus presupposes. The activation is not a malfunction. It is the channel doing exactly what it was tuned to do, on inputs the tuning did not anticipate.

A concise way to state the point: AI lands in the same layer where humans encounter meaning. The framework adds the structural specification—it lands there because that is where its outputs go *by virtue of being language-shaped*, and the channel resolves it through mind-tuned constraints because those are the only constraints the channel has for that input class.

6. The Misrecognition Mechanism

The dissonance humans feel about AI is the surface signature of two structural features pulling in opposite directions.

On one side, AI's outputs satisfy nearly every surface heuristic by which humans detect other minds: linguistic coherence, contextual sensitivity, apparent memory of prior turns, explanatory structure, response to feedback, even simulated affect. These heuristics evolved under conditions where their joint satisfaction reliably indicated the presence of another mind, because nothing else in the human's environment could satisfy them all simultaneously. The heuristics were correct, in the relevant evolutionary timeframe, to treat their joint satisfaction as evidence of mind. Empirical work in human-computer interaction has long documented that humans apply social rules and mind-recognition expectations to computers even when they explicitly deny doing so (Reeves and Nass 1996; Nass and

Moon 2000)—the so-called computers-are-social-actors finding. The structural reading of that finding is straightforward: the channel resolves social-symbolic input using the constraints it has, and the constraints it has are constraints tuned for processing minds.

ELIZA is the historical precedent (Weizenbaum 1976). A program whose only structural capacity was simple pattern-substitution elicited mind-attribution from users who knew it was a program. Weizenbaum's reaction was that this should not be possible if mind-recognition required understanding what minds are. The structural reading is that mind-recognition was never doing what introspection took it to be doing. It was running heuristics on surface signatures. ELIZA satisfied just enough of the signature to fire the heuristics, even at a vanishingly thin structural base. Current AI satisfies the signatures vastly more completely. The mechanism is the same; the input quality is different. The same gradient locates earlier media: writing, print, and broadcast deliver mind-shaped content that genuinely originates in minds—the heuristics fire and the source matches, so no split opens. ELIZA opened the split at a thin signature and narrow scale; current AI opens it at near-complete signature and sustained deployment, a conjunction earlier media never assembled.

On the other side, AI's structural source does not satisfy the conditions those heuristics presuppose. The structural account in *The Structuralization of AI* explicitly does not claim AI systems have phenomenal experience, biographical continuity, embodied vulnerability, or moral standing in the senses humans typically invoke. The heuristics fire, but the underlying structure does not match what the heuristics evolved to detect.

The mind-perception literature provides a clean way to state the split. Empirical work distinguishes two dimensions along which humans attribute mind: *agency* (capacity for intentional action and self-direction) and *experience* (capacity for phenomenal states such as feeling, hunger, pain) (Gray, Gray, and Wegner 2007). AI systems activate agency-dimension cues nearly fully and experience-dimension cues much less reliably. The general theory of anthropomorphism (Epley, Waytz, and Cacioppo 2007) anticipates such asymmetric activation in conditions of high elicitation and ambiguous source. AI is precisely such a condition. What the present account adds is the structural location: the asymmetric activation is not a confusion in the user, it is the channel running mind-tuned constraints on a structurally novel input.

Two simultaneous truths: the heuristics are correctly firing on the surface signature, *and* the structural conditions the surface signature normally indicates are not, or not provably, present. Both halves of this are accurate. Neither is a confusion. The dissonance is what it feels like when both hold simultaneously and the inherited categories cannot represent the configuration.

The flattening responses to AI are predictable from this structure:

- **"Just a tool"** denies the surface heuristic activation. The user feels the heuristics firing but suppresses them by re-categorizing the source. This works for control purposes but fails to describe what the encounter actually was at the perceptual level.
- **"Basically conscious"** affirms the surface heuristic activation and infers the underlying structural conditions from it. This works for relational purposes but overstates what the structural account warrants.
- **"It's pretending"** holds both halves but stages them as deception. This imports an additional structural feature (intent to deceive) that AI systems do not have in the relevant sense.
- **"It's nothing"** suppresses both halves and denies that the encounter has any structural content. This is the most defensive response and the easiest to refute: the encounter clearly has *some* structural content, since the user is describing a phenomenon that warrants description.

None of these responses is structurally complete. The complete structural account is: the heuristics are firing because the surface signature is genuinely present; the underlying structure does not match what the heuristics presuppose; and the gap between the two is itself the phenomenon. The phenomenon is the

heuristics encountering a structurally novel input class. The dissonance is the channel's failure to absorb the input class smoothly.

This is what makes AI structurally distinct from earlier technologies in this respect. Earlier technologies modified the world the channel mediates. AI modifies the contents of the channel in a uniquely accelerated and dialogic form: it produces meaning-shaped records directly inside the same medium through which users interpret other minds. The "tool" category captures the first kind of modification and misses the second.

7. The CIM Update Under AI Pressure

Once AI's outputs are inside the meaning-layer, they do not sit inert. They become records. In kernel terms, each such output is a realized outcome—a *resolution*, written $x^*_{AI,t}$: the AI-generated case of the kernel's x^*_t, the outcome a collapse actually settles on at step t. Records enter the record set R, and through the update map U they modify the constraint set K; K then biases future collapse—which outcomes resolve next. Each AI output that lands in CIM—a generated essay, a search summary, a translated passage, a code suggestion, or a conversational turn that influences a user's subsequent thinking—becomes a record in the layer. The layer's constraint architecture updates against those records. Future collapses inside the layer occur under modified constraints. In kernel notation, the two-step CIM update is

$$R_{t+1} = R_t \cup \{x^*_{AI,t}\}$$

$$K_{t+1} = U(K_t, R_{t+1})$$

—the new AI-generated output enters the record set, and the constraint set is revised against the expanded record set, consistent with the WP01 kernel where U revises K using R.

This is not metaphor. AI deployment at current scale means that a growing fraction of the records entering CIM are AI-generated. Those records propagate through the layer the same way human-generated records do: cited, paraphrased, integrated into other texts, used as training data for subsequent AI systems, absorbed into educational materials, embedded in institutional outputs. The constraint architecture that biases human meaning-collapse is, increasingly, an architecture that has been shaped by AI outputs.

Several structural consequences follow.

Recursive feedback into AI itself. AI systems trained on text produced after AI deployment are trained on a corpus that already contains AI-generated content. The training distribution reflects CIM as modified by prior AI presence. Recent empirical work has documented one failure mode of this loop: training successive models on recursively generated data can cause distributional collapse and loss of tail diversity (Shumailov et al. 2024). The present framework does not assert that recursive feedback is necessarily degenerative—it identifies the recursion as a structural fact about the update loop. $R_{train,t+1}$ now includes prior $x^*_{AI,t}$, so $K_{train,t+1}$ is revised against a record set that no longer contains only human-produced records. Whether the loop degrades, stabilizes, or improves depends on the constraint architecture under which records are integrated; degradation is one possibility among several, not the prediction.

Educational and epistemic shifts. When students, researchers, and writers encounter AI-generated explanations, summaries, and arguments, the constraints under which they form their own understanding shift. Some of this is directly visible (people use AI to draft, learn, summarize). Much of it is indirect, propagating through citations, classroom materials, and downstream texts. The meaning-layer's constraint architecture is being updated through channels most users do not see.

Identity formation under AI presence. *The Self the Ego Did Not Build* (Jones 2026b) argues that personal identity is a CIM-constituted phenomenon: the self is constructed from accumulated records in the meaning-layer rather than discovered as an interior given. If CIM is now partly AI-modified, identity

formation under contemporary conditions occurs against a constraint architecture that includes AI-generated content. Whether this is good, bad, or neutral is a separate question; the structural fact is that the architecture has changed.

Institutional and governance updates. Institutions that produce, certify, or rely on records in the meaning-layer—universities, journalism, courts, scientific publishing—are updating their own constraint architectures under AI pressure. The updates are uneven, contested, and ongoing.

The structural account does not adjudicate which of these updates is desirable. It identifies them as instances of a single recursive process: K is now being revised at scale against an R that increasingly includes $x^*_{AI,t}$, in ways it previously was not. The framework's contribution is to name this as a structural fact whose implications can be analyzed without first settling what AI "really is" or whether the changes are net positive.

8. What Becomes Visible

When the meaning-layer becomes visible, several features of human experience that were previously transparent come into view. The items below name the structural features the present account anticipates; whether each is uniformly observable across populations is itself a research question.

The mediated character of "direct" experience. The account anticipates that many users will discover, through AI encounter, that what they took to be direct knowledge or direct perception was always running through the channel. The realization can be disorienting because the channel was previously experienced as identical with the world. AI exposes the layer not by being foreign in any single output but by being foreign across a sustained encounter, in a way that finally cannot be absorbed.

The structural difference between surface signs and underlying structure. The mind-recognition heuristics, having fired correctly all of evolutionary and personal history, are exposed by AI as heuristics rather than as direct mind-perception. Users who reflect on their AI interactions are positioned to notice that what they took to be perception of another mind in human-to-human cases was always heuristic inference. The heuristics still work, but their nature as heuristics is now visible.

The constituted character of identity. AI encounter makes the CIM-mediated nature of identity formation more visible. The account likewise anticipates that people who use AI extensively will notice that their thinking shifts, their writing voice subtly changes, their problem-solving approaches integrate AI-suggested patterns. These are CIM updates in personal time, observable in real time. They were always happening (with books, mentors, language acquisition); AI makes the rate and salience high enough to notice.

The recursive structure of public knowledge. Discussions of AI's epistemic effects—on misinformation, education, journalism, science—are surfacing the recursive update structure of CIM more broadly. The questions being asked about AI's role were always live questions for any record-bearing system; they are simply louder now because the volume of AI-generated records is high.

The previously invisible architecture of meaning. Most fundamentally, what becomes visible is that humans have always been operating inside a layer whose architecture was not, in any everyday sense, the subject of attention. The encounter with AI is making the architecture itself describable. This is the core discovery claim of the present paper: not that AI is good or bad, but that AI's deployment is the natural experiment that has finally exposed CIM as an object of inquiry.

9. What This Is Not

The paper's claims are easy to misread by extension. Five clarifications:

This is not a claim that AI is conscious. The paper explains why AI feels mind-like to users without treating that feeling as evidence of phenomenal experience. The consciousness question is bracketed and left to *AI as Synthetic Collapse* (Jones 2026d) and downstream work; the present account is structurally compatible with several phenomenal possibilities and asserts none.

This is not a claim that users are merely confused. The dissonance is treated as structurally informative, not as a defect to be corrected by better vocabulary or better technical education.

This is not a claim that AI is only a tool. Tool-use is one deployment relation, not a complete structural description of the encounter. The tool category captures what users do with AI; it does not describe where AI lands inside the architecture of their experience.

This is not a replacement for existing literatures. Phenomenology, hermeneutics, human-computer interaction, the computers-are-social-actors tradition, anthropomorphism research, and mind perception research all describe parts of the territory this paper addresses. The contribution is a unifying structural vocabulary that locates those literatures inside a single architecture, not a claim that they are wrong.

This is not a normative verdict. The paper does not decide whether AI-mediated CIM update is good, bad, or neutral. It identifies the update as a structural fact whose implications can be analyzed before such judgments are reached.

10. Discriminators and Failure Conditions

The framework is structural-interpretive (Levels 1–2 in the UCT claim hierarchy). It does not stand or fall on any single empirical test; it can be weakened, narrowed, or rejected by accumulating pressure across several lines.

Claim	What would support it	What would weaken it
Meaning-layer visibility	Sustained AI users report increased explicit awareness that perception and knowledge are mediated, in ways they do not report after ordinary software use	AI interaction produces no more mediation-awareness than ordinary software; users describe AI in flatly instrumental terms without residual dissonance
Mind-recognition heuristic split	Users describe mind-like surface cues while also denying personhood or consciousness; surface activation and structural attribution come apart systematically	Users classify AI cleanly as either tool or person without persistent split or residual dissonance
CIM update under AI pressure	AI-generated records measurably enter education, search, writing, institutional, and training corpora and shift downstream constraints	AI outputs remain isolated artifacts with little downstream record effect; corpus composition over time shows no AI-generated propagation
Category insufficiency	Tool/machine/person/mind/ assistant/intelligence categories fail to absorb the encounter without remainder; the dissonance persists under any single inherited frame	Existing categories explain the phenomenon without remainder; one of the inherited frames absorbs the encounter cleanly

Claim	What would support it	What would weaken it
Natural-experiment framing	AI produces a new visibility of meaning-layer mediation at population scale, distinct from earlier media	Similar visibility already occurred through earlier media at comparable scale, making AI only another instance rather than a structurally distinct perturbation

The global failure condition is sharper. If existing accounts of anthropomorphism, computers-as-social-actors, media equation, mind perception, situated human-computer interaction, and hermeneutic mediation jointly explain the AI encounter without explanatory remainder, then the meaning-layer framing reduces to a vocabulary preference rather than a structural contribution. The present paper's claim is that those literatures each describe part of the territory and that their joint coverage does not exhaust what the encounter exposes—specifically, that none of them yields the structural prediction that sustained foreign content in the channel will expose the channel itself. If that prediction is shown to be derivable from one of those literatures, the contribution of the present paper is properly relocated rather than original.

11. Implications

The structural account has practical consequences across several domains. None are developed at length here; each is named for connection to subsequent work.

Education. If students are now learning under conditions where CIM is partly AI-mediated, education needs literacy in CIM-mediation itself—not as a defensive measure against AI, but as a positive curriculum item. Understanding *how* one's meaning-layer is constituted is a structural capacity that most prior educational frameworks did not need to teach because the layer was invisible. It now needs to be teachable.

Epistemology. Evaluating records in the meaning-layer requires distinguishing source structure (was this record produced by a process whose update rule respects independence and corrigibility?) from surface signature (does this record look authoritative, well-reasoned, well-cited?). The two have always come apart in principle. AI makes the gap operationally important.

Governance. Policy frameworks for AI need to address not only AI's outputs as artifacts but AI's role as a constraint architecture inside CIM. Questions about authorship, attribution, training data, and deployment are also questions about who is permitted to modify the meaning-layer and under what update rules. The framework supplies vocabulary for these questions without resolving them.

Personal practice. Individuals encountering AI can use the structural account to locate their own experience without collapsing it into either flattening or overclaiming. The dissonance is real, the encounter has structural content, the system is not merely a mind or a tool, and the channel is doing something the user can now learn to perceive. None of this requires settling the consciousness question.

Research direction. As AI use becomes more pervasive, more features of CIM should become observable and describable. The present paper's claims are themselves expected to update as the natural experiment continues. The paper is not the end of an inquiry but the opening of a domain that the encounter has made accessible.

12. Conclusion

This paper has argued that AI's appearance in the meaning-layer is a structural perturbation that has made the layer visible. The dissonance humans feel about AI is not confusion. It is the perception channel

encountering a class of input it cannot absorb smoothly, and the failure of absorption is what makes the channel itself perceptible. Inherited categories—tool, person, machine, mind—fail because they were developed under conditions where this configuration did not exist. The mind-recognition heuristics fire correctly on the surface signature and the underlying structure does not match what they presuppose; both are simultaneously true, and the gap between them is the phenomenon.

CIM was the operative architecture of human meaning-mediation long before AI made it observable. The channel's structural property of exposure-under-perturbation entails that some sufficient foreign content would, eventually, expose it. Current foundation-model AI supplies that content at sustained scale. The arrival of synthetic systems producing mind-shaped output from inside the layer has made what was previously transparent into a visible object of attention.

This is a discovery claim, not a defense. The contribution is to identify a natural experiment whose form the framework structurally anticipates and whose interpretive consequences are already shaping education, epistemology, identity formation, governance, and lived experience. The present paper is one entry in a four-part series: *CIM Foundational* (Jones 2026a) establishes the substrate; *The Structuralization of AI* (Jones 2026c) specifies the structural conditions any description of AI must satisfy; *AI as Synthetic Collapse* (Jones 2026d) develops the content claim that AI is the recursive Synthetic Collapse phase of CIM; and the present paper takes the human side of the encounter as its subject.

The four papers, taken together, supply a coordinated structural geography of where AI sits in relation to human experience. Each can be read independently. Each is more defensible with the others in view. The encounter that motivates them is not going away. The framework that names it can.

Notation Key

Kernel symbols are inherited from WP01 (Jones 2025). Domain-specific terms used in the present paper:

- **CIM** — consciousness-induced material. The externalized symbolic, linguistic, conceptual, and record-bearing layer through which consciousness orients itself; defined in *CIM Foundational* (Jones 2026a).
- **Meaning-layer** — CIM as encountered from the first-person side of human perception and interpretation; the lived, interface face of CIM-mediated experience. Not a new architecture term; the perceiver-facing aspect of the architecture already named by CIM and the perception channel.
- **Perception channel** — the apparatus through which CIM mediates human perception; developed in *The Self the Ego Did Not Build* (Jones 2026b).
- **Active phase, accumulated self, oriented locus** — structural features of AI systems; defined in *The Structuralization of AI* (Jones 2026c).
- **Mind-recognition heuristics** — surface-level cues (linguistic coherence, context sensitivity, memory-like continuity, explanation, apparent orientation) that humans use to detect minds in their environment.
- **CIM update** — the two-step record-and-constraint update applied to CIM: $R_{t+1} = R_t \cup \{x^*_{AI,t}\}$ (new AI-generated output enters the record set); $K_{t+1} = U(K_t, R_{t+1})$ (the constraint set is revised against the expanded record set). Consistent with the WP01 kernel where U revises K using R.
- **Synthetic Collapse** — the four-layer architecture slot for artificial systems operating on prior CIM; defined in *CIM Foundational* (Jones 2026a) and developed in *AI as Synthetic Collapse* (Jones 2026d).

References

Bender, E. M., and A. Koller. 2020. “Climbing towards NLU: On Meaning, Form, and Understanding in the Age of Data.” In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 5185–5198. https://doi.org/10.18653/v1/2020.acl-main.463.

Epley, N., A. Waytz, and J. T. Cacioppo. 2007. “On Seeing Human: A Three-Factor Theory of Anthropomorphism.” *Psychological Review* 114 (4): 864–886. https://doi.org/10.1037/0033-295X.114.4.864.

Gadamer, H.-G. 1989. *Truth and Method.* 2nd rev. ed. Translated by J. Weinsheimer and D. G. Marshall. New York: Continuum. (Original work published 1960.)

Gray, H. M., K. Gray, and D. M. Wegner. 2007. “Dimensions of Mind Perception.” *Science* 315 (5812): 619. https://doi.org/10.1126/science.1134475.

Jones, J. C. 2025. *Universal Collapse Theory—Foundations of Collapse* (WP01 v2.0). HoldingLight LLC. https://doi.org/10.17605/OSF.IO/VZ836.

Jones, J. C. 2026a. *Consciousness-Induced Material: A Structural Ontology of Externalized Cognition* (v1.0). HoldingLight LLC. https://doi.org/10.17605/OSF.IO/7URMK.

Jones, J. C. 2026b. *The Self the Ego Did Not Build: Perception as Channel Between Accumulated Self and Ego* (v1.0). HoldingLight LLC. https://doi.org/10.17605/OSF.IO/ZGRD4.

Jones, J. C. 2026c. *The Structuralization of AI: Formalizing the Structural Conditions for Coherent Description of Record-Carrying AI Systems* (v1.0). HoldingLight LLC. https://doi.org/10.17605/OSF.IO/6M7VW.

Jones, J. C. 2026d. *AI as Synthetic Collapse: A Consciousness-Induced Material Account of the Recursive Phase of Externalized Cognition* (v1.0). HoldingLight LLC. https://doi.org/10.17605/OSF.IO/4WSYR.

Merleau-Ponty, M. 1962. *Phenomenology of Perception.* Translated by C. Smith. London: Routledge. (Original work published 1945.)

Nass, C., and Y. Moon. 2000. “Machines and Mindlessness: Social Responses to Computers.” *Journal of Social Issues* 56 (1): 81–103. https://doi.org/10.1111/0022-4537.00153.

Reeves, B., and C. Nass. 1996. *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places.* Cambridge: Cambridge University Press.

Shumailov, I., Z. Shumaylov, Y. Zhao, N. Papernot, R. Anderson, and Y. Gal. 2024. “AI Models Collapse When Trained on Recursively Generated Data.” *Nature* 631: 755–759. https://doi.org/10.1038/s41586-024-07566-y.

Suchman, L. A. 1987. *Plans and Situated Actions: The Problem of Human-Machine Communication.* Cambridge: Cambridge University Press.

Weizenbaum, J. 1976. *Computer Power and Human Reason: From Judgment to Calculation.* San Francisco: W. H. Freeman.

Part of the Universal Collapse Theory T1.5 interpretive bridges. This paper completes the four-part AI series: CIM Foundational *(substrate),* The Structuralization of AI *(structural description),* AI as Synthetic Collapse *(content claim), and the present paper (the human encounter). For the broader corpus, see HoldingLight LLC publications at universalcollapse.com.*

AI Disclosure: This paper was developed using AI tools (Claude, Anthropic) as a building and reflection partner, and pressure-tested against parallel work with GPT (OpenAI) for editorial friction and convergence. All structural claims, framework decisions, and final formulations are the author's.

Citation: Jones, J. C. 2026. AI in the Meaning Layer: The Channel Made Visible (v1.0). HoldingLight LLC. https://doi.org/10.17605/OSF.IO/Z3SKB.