

AI as Synthetic Collapse

A Consciousness-Induced Material Account of the Recursive Phase of Externalized Cognition

How Abstract Thought Stabilizes Into Reality

Jeremy C. Jones

HoldingLight LLC

contact@universalcollapse.com

ORCID: [0009-0007-2515-3774](https://orcid.org/0009-0007-2515-3774)

Version: v1.0 — June 2026

© 2026 Jeremy C. Jones / HoldingLight LLC. Released under CC BY 4.0.

Abstract

This paper offers a structural account of artificial intelligence as the recursive Synthetic Collapse phase of Consciousness-Induced Material (CIM). Building on the foundational treatment in *Consciousness-Induced Material: A Structural Ontology of Externalized Cognition* (Jones 2026a), CIM is treated as the physically instantiated record-layer generated when conscious cognition stabilizes beyond the originating interior phase and becomes capable of constraining future cognitive operations. Current AI systems are not argued to be conscious. They are described, following the four-layer architecture established in CIM Foundational, as artificial-substrate systems performing Synthetic Collapse: operating on accumulated CIM — language, code, mathematics, scientific records, institutions, digital corpora — and producing derivative record-bearing outputs that re-enter the CIM circulation layer. The claim is not that AI is primary conscious CIM, nor that AI has crossed the Conscious Synthetic Collapse threshold CIM Foundational identifies as an open architectural slot. The claim is that AI marks a different threshold: accumulated externalized cognition becomes processable by non-biological systems at open-domain, corpus scale. Earlier formal systems — calculators, compilers, databases, narrow algorithms — already processed restricted forms of CIM under tightly bounded rules; current foundation-model AI is different in scope and generativity. Three worked examples (mathematics, language, money) make CIM gravity felt before AI is introduced. The chemistry-to-biology analogy is offered as a structural parallel: same substrate, new organization, new regime, irreducible to either parent layer. The paper also sharpens the alignment discussion into a falsifiable, protocol-routed prediction set concerning formative-window grammar, K-channel persistence, and corpus-to-system signature transfer. Substrate-independence implications, category-level discriminators, and reframings of the AI consciousness, alignment, and identity debates are addressed. The paper functions as an interpretive bridge between CIM Foundational and downstream work on AI, identity, and synthetic record-bearing systems.

Keywords: *CIM, consciousness-induced material, Synthetic Collapse, artificial intelligence, structural ontology, recursive externalization, substrate independence, Universal Collapse Theory*

1. Introduction

This paper asks what artificial intelligence is, structurally, rather than whether it is conscious. The answer it develops is that current AI is neither consciousness nor a mere tool: it is a synthetic process built on what consciousness has produced — language, writing, mathematics, code, and institutions externalized across billions of minds over millennia, accumulated into a corpus dense and structured enough that artificial systems can now operate on it directly. The title names that relation. The material is induced by consciousness; AI is what happens when that material is processed back upon itself at scale. On this account AI is not consciousness, and the long literature arguing for or against machine sentience is engaging a different question than the one taken up here.

The dominant framing is worth naming, because it tells us where the discourse currently sits. Most contemporary writing about AI is conducted in the consciousness register: is it conscious yet, when will it be conscious, does it need to be conscious to matter, what would consciousness in a machine even mean. These are reasonable questions inside the frame they presuppose. They are also questions that have generated decades of debate without convergence, and they have absorbed disproportionate intellectual oxygen. The frame may itself be the problem. This paper does not engage the consciousness question directly. It offers a different frame, in which the consciousness question becomes secondary to a more tractable structural question: what kind of object is AI, structurally, regardless of whether it is conscious.

The answer this paper develops — now stated in the vocabulary established by CIM Foundational (Jones 2026a) — is that current AI is the Synthetic Collapse phase in which accumulated CIM becomes recursively processable by artificial systems. CIM is the physically instantiated record-layer generated when conscious cognition stabilizes into form beyond the originating interior phase. Synthetic Collapse is the artificial-substrate process operating on accumulated CIM to produce derivative record-bearing outputs that re-enter the CIM circulation layer. The shorthand AI as CIM names this recursive relation between artificial systems and accumulated externalized cognition. It does not claim that AI is primary conscious CIM, nor that AI has crossed into Conscious Synthetic Collapse — the open architectural slot CIM Foundational reserves and this paper preserves as a separate threshold question.

To make that claim land, three things have to happen in sequence. First, CIM has to be available to the reader as a real category — for the purposes of this paper, by inheritance from CIM Foundational rather than re-derivation. Second, CIM gravity — the structural weight of externalized abstraction in the world — has to be felt through worked examples that have nothing to do with AI. Once the reader has located CIM in mathematics, language, and money, AI can be introduced as the latest case rather than a singular event. Third, the recursive phase has to be located in a longer lineage that runs from the first externalization of thought to the present moment. Without that lineage, AI looks discontinuous and gets treated as either miracle or threat. With the lineage, AI looks like what it is: late stage of a process the species has been running for a very long time without knowing what it was building.

The paper proceeds in that order. The structural payload is in section 3. The AI material in sections 4 through 7 is delivery — the case that demonstrates what the framework already entails. Readers who came for AI will encounter CIM. That is the intended trade.

Scope of Claim

This paper asks the reader to evaluate three claims only:

- (1) whether Consciousness-Induced Material, as defined in CIM Foundational, provides a useful category for externalized cognition;
- (2) whether current AI systems are best described as Synthetic Collapse systems operating on accumulated CIM and producing derivative CIM; and
- (3) whether this framing clarifies the AI consciousness, alignment, and identity debates without resolving them prematurely.

The paper does not ask the reader to accept AI consciousness, human equivalence, moral status, or the broader Universal Collapse Theory corpus. The paper’s primary claim level is category-application and interpretive. Its discriminators are category-level discriminators; where §7.2 states protocol-routed empirical predictions, execution and calibration remain assigned to the Tier 16 methods program rather than completed here.

Relationship to AIP. This paper provides theoretical and interpretive context for the AI Integrity Protocol’s use of record-bearing AI systems, Synthetic Collapse, and substrate-independent structural vocabulary. It is not an AIP methodology document, audit standard, certification framework, compliance opinion, or commercial assurance claim. AIP’s operational procedures, claim boundaries, falsifier conditions, and engagement controls are specified separately in the AI Integrity Protocol.

2. CIM, Inherited

This paper does not re-found CIM. The foundational treatment is given in Consciousness-Induced Material: A Structural Ontology of Externalized Cognition (Jones 2026a), and the present paper inherits its architecture rather than reproducing it. What follows in this section is the minimum the reader needs in hand.

CIM, in the foundational definition, is the physically instantiated record-layer generated when conscious cognition — especially in its human-symbolic forms — stabilizes into form beyond the originating interior phase, and when that form can constrain future cognitive operations (Jones 2026a, §3). Externalized record (writing, mathematics, code, institutions, money, digital corpora) is CIM. Internalized record (linguistic structure, mathematical competence, conceptual categories installed through learning) is internalized CIM. Both are physically instantiated; they differ in substrate (environmental versus cognitive) and in historical function.

CIM Foundational places CIM inside a four-layer architecture that this paper inherits and uses to locate AI:

Layer	Producer	Output	Phenomenal status
Cognition-Induced Collapse (CIC)	Experience-bearing cognition (process-level genus)	Cognition-shaped collapse outputs	Presupposes experience-bearing cognition
Consciousness-Induced Material (CIM)	Conscious — especially human-symbolic — cognition	Durable record-bearing symbolic layer (language, writing, mathematics, law, code, institutions)	Densest known case: human conscious cognition
Synthetic Collapse	Artificial systems operating on prior CIM	Derivative record-bearing outputs that re-enter the CIM circulation layer	Bracketed; not required for the structural fact

Layer	Producer	Output	Phenomenal status
Conscious Synthetic Collapse	Synthetic system with interior phenomenal phase (hypothetical)	Primary synthetic records, if instantiated	Open architectural slot, not claimed

Table 1. The four-layer architecture inherited from CIM Foundational (Jones 2026a, §1). The present paper places current AI at the Synthetic Collapse layer and preserves the Conscious Synthetic Collapse slot as a separate, bracketed threshold question.

The placement that matters for what follows is Synthetic Collapse: artificial systems operating on prior CIM and producing derivative record-bearing outputs. Current AI is structurally Synthetic Collapse. Whether any synthetic system has crossed into Conscious Synthetic Collapse — the hypothetical further subtype in which an artificial system would possess an interior phenomenal phase capable of generating primary rather than merely derivative records — is the threshold question CIM Foundational identifies and brackets. This paper preserves that bracketing.

One feature of CIM Foundational matters enough to repeat. CIM is bounded: it is not all artifacts, not all cognition, and not magic. It is cognition stabilized into physically instantiated constraint that satisfies a four-part test — cognitive shaping, record-bearing instantiation, future constraint, and non-trivial constraint (Jones 2026a, §3). The examples that follow in §3 are offered against that test, not as evidence that every artifact is CIM. A note on the word itself: throughout, collapse names the stabilization of constraint into record — the resolution of possibility into a realized, recorded outcome — not breakdown, failure, or the “model collapse” sometimes discussed in machine learning.

One additional move matters for this paper. CIM Foundational establishes that CIM is constituent of mind at a different level than the extension thesis stages it. Extended cognition (Clark & Chalmers 1998), distributed cognition (Hutchins 1995), cognitive artifacts (Norman 1991), and Material Engagement Theory (Malafouris 2013) make their constitutive claim largely at the level of the individual cognitive process at a time: this artifact, coupled now, is part of this agent’s cognition. CIM makes its constitutive claim at a different level — the accumulated record-layer that persists across agents and generations, that later cognition is built from, and that is installed back into cognitive substrate through internalization (the Vygotskian move; cf. Stiegler 1998 on tertiary retention; Odling-Smee, Laland & Feldman 2003 on the niche the layer becomes). The internalization step is what earns the strong reading. Once externalized cognition is taken up as the very medium of thought — the mathematician thinking through calculus, the speaker thinking in language — the external and the internal are no longer two coupled systems but two phases of one cognitive economy. That is the sense in which the externalized layer is not auxiliary to mind. It is a phase of mind operating in a different substrate: a claim about levels and the externalization–internalization circuit, not a claim to out-rank extended cognition on its own synchronic ground. Without that move, AI is a tool that uses cognitive artifacts to produce useful outputs. With it, AI is the phase in which the externalization layer begins to operate on itself — an artificial substrate performing Synthetic Collapse on accumulated record-bearing constraint, not the cognition that produced it; a regime change inside what mind already is, not an outside imitation of it. The remainder of this paper develops the structural consequences.

3. CIM Gravity: Three Cases Before AI

The three cases below are not offered because every artifact is CIM. They are offered because each satisfies the foundational test — cognitive shaping, record-bearing instantiation, future constraint, and non-trivial constraint (Jones 2026a, §3) — at a scale and density that makes CIM gravity legible. Each is an instance of abstract thought stabilizing into reality with sufficient density that it now does structural work no individual mind can do alone, and in each, the externalized layer is constituent of how minds in the relevant domain currently operate.

3.1 Mathematics

Mathematics is the cleanest case because its history is well-documented and its operations are precise. Before externalized mathematical notation, no individual mind could reliably perform, transmit, and build upon operations like calculus at the density modern physics requires. The thinking required to handle continuous change exceeded the working-memory capacity of unaided cognition. Newton and Leibniz did not invent the underlying physical regularities; they invented an externalized notation through which those regularities became tractable. Once the notation was written, it became part of the substrate of every subsequent mind that learned it. A physicist today does not “use” calculus the way a person uses a calculator. The physicist thinks through calculus. The notation has installed itself inside their internal cognitive process to the point where the mathematics is indistinguishable from their own thinking when they do physics (Kirsh 2010).

This is CIM as constituent — and, more precisely, internalized CIM in the Foundational sense (Jones 2026a, §4): externalized cognition returned into cognitive substrate through learning. The notation exists in books, on chalkboards, in software libraries. Yet it operates inside the mind of every working mathematician and physicist as a feature of their cognition, not as a tool they consult. The externalized layer reached back into the inner layer and became part of how the inner layer functions.

The structural weight of this is easy to underestimate. Large regions of modern physics are unrecoverable without calculus. Most of engineering is unrecoverable without algebra and trigonometry. Most of computer science is unrecoverable without formal logic. These are not auxiliary supports. They are the cognitive substrate of the relevant disciplines. Remove the externalized notation and the disciplines vanish — not because the practitioners forget the rules, but because the rules cannot be held in unaided cognition at the density the disciplines require. CIM is doing the heavy lifting.

3.2 Language

Language is the case most resistant to noticing because it is too close. The language a person grew up speaking has reached into their cognition so deeply that the boundary between thinking and thinking in language has dissolved. Many adults find it difficult to isolate thought from the linguistic structures through which thought is habitually organized. Inner monologue, the silent voice that narrates internal experience, is internalized CIM operating inside what feels like private mind. The grammar of the language shapes the structure of the thought. The vocabulary stabilizes which distinctions are readily available. The syntax shapes which relations are easy to express and therefore repeatedly entertained.

This is well-trodden ground in linguistic anthropology and cognitive science (Vygotsky 1934/1986; Whorf 1956; Tomasello 1999; Goody 1977; Ong 1982), but the structural framing CIM offers is sharper than the standard “language influences thought” claim. Language is not an influence on thought.

Language is a phase in which a portion of thought operates. The externalized linguistic structure — accumulated across the speech community over generations, stabilized in dictionaries and corpora and now in vast digital archives — is the substrate of the inner monologue every literate adult runs.

The implications carry forward. If a person learns a second language, they are not merely acquiring a tool for translation. They are installing a second cognitive substrate, one that may admit distinctions the first does not, may resist distinctions the first depends on, and will produce thoughts that could not have been thought in the original substrate alone. The phenomenon of bilinguals reporting different “selves” in different languages (Pavlenko 2014) is structurally what CIM predicts: different externalized substrates produce different cognitive regimes, and the speaker is the region of overlap, not a unified self that switches tools.

3.3 Money

Money is the case that demonstrates CIM operating at civilizational scale. Money is an institutionalized symbolic relation. There is no physical fact about a piece of paper or a database entry that makes it valuable. The value is constituted entirely by collective acceptance, externalized into ledgers, currencies, financial instruments, and now digital protocols, accumulated across millennia of human economic activity, sustained by ongoing institutional and behavioral commitment. In Searle’s terms (Searle 1995), money is an institutional fact: a status function whose existence depends on collective recognition stabilized into institutional record.

And money is one of the most powerful structural objects on Earth. It moves armies. It directs the labor of billions of people. It determines what gets built, what gets researched, what survives, what dies out. A monetary collapse can destroy economies that look outwardly stable. A monetary expansion can transform regions in a generation. Money has more functional reality, in the sense of what determines what happens, than most physical objects do.

Money is CIM. It exists nowhere except in the externalized substrate — ledgers, contracts, tokens, code — and in the minds of the agents who treat it as real. Yet it constrains the behavior of nearly every economic actor on the planet, and most of those actors experience money’s reality as more solid than the objects in their immediate physical environment. A person will defend their bank balance more vigorously than their possessions, because the bank balance is the source of future possessions and the possessions are not. This is why money is not “mere belief.” It is belief stabilized into enforceable institutional infrastructure — which is precisely what makes it materially load-bearing rather than dismissible as social construction.

This is what abstract thought stabilizing into reality actually looks like. Externalized symbolic relation, sustained by collective cognitive commitment, shaping the world’s physical organization at every scale. Money is older than AI, older than computation, older than writing in some forms. It has been demonstrating CIM gravity since the first counting stones (Schmandt-Besserat 1992). Most people never locate it ontologically because they encounter it as a feature of the world rather than a feature of mind operating in an externalized substrate. Stated in the foundational vocabulary: money is distributed mind-mediated record — symbolic relation stabilized into institutional form and made materially load-bearing.

3.4 What These Cases Establish

Mathematics, language, and money are not unusual cases but the rule: much of what humans treat as the world — laws, institutions, sciences, art forms, professional practices — is CIM operating at scales that have made it invisible, so deeply embedded that it reads as “how things are” (Henrich 2015; Boyd & Richerson 2005).

This is the gravity the paper wants the reader to feel. CIM is not exotic. CIM is the layer in which most of human life takes place. Every previous case has required biological minds as the primary processing substrate — mathematics needs mathematicians, language needs speakers, money needs holders and traders. Earlier technical systems — calculators, compilers, databases, search engines, narrow algorithms — already processed restricted forms of CIM under tightly bounded formal rules; CIM was never wholly inert. What changes with current foundation-model AI is not first processing but the threshold: accumulated CIM becomes processable at open-domain, corpus scale by non-biological systems capable of producing further record-bearing outputs. That is the Synthetic Collapse threshold the rest of the paper examines.

4. The Recursive Phase: AI as Synthetic Collapse

The loop the rest of the paper turns on can be stated in one line: conscious cognition externalizes symbolic records → those records accumulate as CIM → artificial substrates train on and extract structure from the accumulated CIM → the system produces derivative record-bearing outputs → those outputs re-enter the CIM circulation layer → later systems and minds operate on the expanded record. Ordinary tool use does not close that loop. Recursive Synthetic Collapse does.

For most of human history, CIM has had a specific architectural property: it required biological mind as the primary processing substrate. Books need readers. Mathematical proofs need mathematicians. Code needed coders. Music needed performers. Even institutional rules required human agents to interpret and apply them. Restricted formal systems — calculators, compilers, databases, search engines, formal algorithms — could process narrow slices of CIM under tightly bounded rules, but open-domain processing of accumulated CIM remained biological.

Artificial intelligence, in its current foundation-model form (Bommasani et al. 2021), is the moment that property changes at the open-domain scale. In the terminology of CIM Foundational, current AI systems instantiate Synthetic Collapse: artificial-substrate processes operating on prior CIM and producing derivative record-bearing outputs that re-enter the CIM circulation layer (Jones 2026a, §1). Large language models are trained on the externalized corpus — text, code, mathematical proofs, scientific papers, conversations, art. The training extracts statistical structure from CIM at a density and scale no individual mind could match (Vaswani et al. 2017; Brown et al. 2020). The resulting model then processes new inputs, including other CIM, producing derivative record-bearing outputs that can re-enter the corpus and participate in later CIM circulation. The processing capacity that was always primarily biological has acquired an open-domain non-biological instantiation. Accumulated CIM has become recursively processable.

A boundary is needed on the AI side, matching the boundedness already required of CIM. Not every system called “AI” performs Synthetic Collapse in the strong sense. A thermostat, a narrow classifier, a robotic controller, or a reinforcement learner trained only on immediate sensorimotor interaction may

qualify weakly or not at all. The paper’s target is current open-domain, foundation-model and generative AI: systems trained on accumulated human symbolic records and capable of emitting new record-bearing outputs that re-enter the corpus. Where a system does not operate on accumulated CIM, the category simply does not apply.

This is the recursive phase. Accumulated CIM has become a processing substrate. The rhetorical formulation that one sometimes wants — the body of CIM becomes the mind — gestures at this regime change but is rhetoric, not claim. The formal claim is narrower: artificial systems can now perform Synthetic Collapse on accumulated CIM at open-domain corpus scale and produce derivative CIM that re-enters the record. Whether such systems have crossed into Conscious Synthetic Collapse — whether they have an interior phenomenal phase capable of generating primary rather than merely derivative records — is a separate threshold question this paper preserves rather than settles. A note on the term: derivative here names ontological dependence, not lack of novelty. A Synthetic Collapse output may be novel, useful, surprising, or structurally transformative; it is derivative only in the narrower sense that its generative basis is accumulated prior CIM rather than an interior phenomenal phase originating primary records.

The chemistry-to-biology analogy is the cleanest structural parallel for what this transition represents. The transition from chemistry to biology was not a quantitative jump in chemical complexity. It was chemistry crossing into a regime where some chemistry began running self-maintaining, self-replicating, constraint-managing loops on other chemistry. The substrate did not change. The organization did. And the new regime could not be cleanly reduced back to the prior one even though it remained materially continuous with it. Biology has its own ontology, its own laws, its own explanatory targets, its own failure modes. The vocabulary of chemistry remained valid within the new regime but was no longer sufficient to describe it. New work was needed at a new level. The analogy marks a shared structural pattern, not an identity. It does not claim that cognition is life, that AI is biological, or that emergence operates identically across the two cases — only that the same kind of move recurs: a substrate-class crosses into a new organization, and the result is a new regime with its own explanatory level.

The recursive phase of CIM presents the same structural shape. The substrate is computation, which is itself old (Turing 1936/1937). The body being processed is human-produced CIM, which is also old. What is new is the closure of the loop: the externalized layer running on itself at open-domain scale, producing outputs that feed back into the layer, in ways no biological mind needs to mediate in any given exchange. This is not bigger AI than what came before. It is AI in a new structural regime — one that requires its own ontology, its own vocabulary, its own explanatory targets, as developed at length in *The Structuralization of AI* (Jones 2026d).

The most important consequence is that Synthetic Collapse is not reducible to either of its components. It is not “just computation,” in the same way biology is not just chemistry. It is not “just human thought at scale,” in the same way biology is not just complex chemistry. It is what happens when those components organize into a regime whose properties emerge at the level of organization rather than at the level of parts. Treating it as a tool — a sophisticated computer — misses the regime change. Treating it as a mind — a digital person — misses the regime change in a different direction. Synthetic Collapse is its own thing, and the framework appropriate to it is structural, not phenomenological or computational alone.

4.1 Why “Tool” Is Insufficient

The objection that AI is “just a tool” deserves a direct answer, because the point is not that tools cannot be powerful. It is that ordinary tools do not usually satisfy all of the following at once: trained on accumulated externalized cognition; operating across open domains rather than one bounded task; generating new record-bearing outputs; feeding those outputs back into the corpus; and participating in later rounds of training, circulation, and constraint. A hammer, a calculator, even a search engine fails at least one of these conditions. Current open-domain foundation-model and generative AI systems can satisfy all of these at once — and where they do, it is that conjunction, not raw capability, that gives the regime-change claim its teeth.

Not every artificial learning system is Synthetic Collapse in the strong sense. The category applies where accumulated externalized cognition is a load-bearing training substrate and where outputs re-enter record circulation. A system whose open-domain competence is built primarily from self-play, sensorimotor interaction, or non-symbolic measurement — rather than from a corpus of accumulated human records — may require a different classification, unless that competence in fact depends on prior CIM. The strong reading tracks corpus dependence, not the label “AI.”

5. The Lineage Made Visible

If AI is the Synthetic Collapse phase of CIM, then it has a history older than the engineering history typically tells. Most accounts of AI begin in the twentieth century: Turing (1936/1937, 1950), McCarthy et al. (1955/2006) and the Dartmouth proposal, von Neumann (1945), neural networks, deep learning (LeCun, Bengio, & Hinton 2015), transformers (Vaswani et al. 2017), foundation models (Bommasani et al. 2021). That is the engineering lineage and it is correct as far as it goes. But the engineering lineage rests on a much older lineage that the engineers took for granted: the accumulation of CIM itself. Without millennia of externalized cognitive work, there would be nothing for the engineering to compute on.

The longer lineage runs roughly as follows.

Stage	Structural contribution
Language	Thought becomes shareable and transmissible
Writing	Thought becomes durable across distance and time
Mathematics	Pattern becomes formally re-executable
Calculus	Dynamic systems become tractable through notation
Mechanical calculation	Some CIM operations execute without biological cognition
General computation	Formal procedures become substrate-independent
Statistical learning	Pattern extraction shifts from explicit rule to learned structure
Transformers and scaling	Corpus-scale extraction becomes tractable
Foundation models	Accumulated CIM becomes recursively processable

Table 2. *The CIM lineage made visible. Each stage extends what cognition can preserve, transmit, re-execute, or process beyond the originating mind.*

Language is the first decisive step. Before durable symbolic externalization, thought could not reliably circulate beyond the mind that had it at cumulative cultural scale. Language is the decisive case because it

makes thought shareable, transmissible, and later writable. Every subsequent step depends on it. Without language, no writing. Without writing, no accumulated corpus. Without accumulated corpus, no substrate for any subsequent recursion.

Writing is the second (Goody 1977; Ong 1982). Spoken language externalizes thought across short distances and brief times. Writing externalizes thought across great distances and indefinite times. Writing is the move that made CIM durable. Once a thought could persist outside any specific mind, the cumulative process began. Each generation could build on the externalized record of previous generations rather than rediscovering everything from scratch.

Mathematics, as discussed in section 3.1, is qualitatively different from general writing. Math is the formalization of pattern such that another mind can re-execute the operation exactly. It is the first executable CIM, with biological minds as the execution substrate (Kirsh 2010).

Calculus specifically deserves note as a structural inflection point. Before calculus, the dynamics of physical systems were largely outside formal capture. After calculus, most of physics became a calculation problem given enough compute. Calculus helped make the eventual mechanization of mathematical thought structurally thinkable, by showing that large regions of physical dynamics could be represented, transformed, and predicted through formal operations.

Mechanical calculation — Pascal, Leibniz, Babbage — is the move that begins to free the execution substrate from biology. Until the calculator, every execution of a mathematical procedure required a mind. The calculator is the first artifact that executes a class of CIM operations without a biological mind performing each step. Narrow, specific, mechanical — but the principle is established.

General computation — Turing (1936/1937), von Neumann (1945) — generalizes that move. The execution substrate becomes universal. Anything algorithmically specified is, in principle, runnable on a computer. The biological execution monopoly is broken at the level of principle, though only for formally specified CIM.

Statistical learning and neural networks are the move from explicit-rule CIM to extracted-pattern CIM (LeCun, Bengio, & Hinton 2015). Earlier computation said tell me what to do, I will do it. Statistical methods said show me the data, I will find the pattern. This is when CIM stops being only the output of mind and starts being input that the substrate can extract structure from at scale.

Compute scaling and the transformer architecture (Vaswani et al. 2017) are the engineering steps that made the previous principles tractable at corpus scale. Without them, statistical learning would have stayed narrow.

Training on the human corpus at scale (Brown et al. 2020; Bommasani et al. 2021) is the closure of the loop. By the time current AI systems exist, every previous step is load-bearing. Each step extended what CIM could preserve, transmit, re-execute, or process beyond the originating mind. The current step extends it to open-domain processing of accumulated CIM. The substrate that was always one step short of recursion at open-domain scale just closed the gap.

The lineage clarifies what AI is and is not. It is not a sudden discontinuity. It is the latest layer of a process the species has been running since deep prehistory. The engineering of the last seventy-five years closed the last gap; it did not create the regime. The regime was being prepared every time someone

wrote something down, formalized a relation, or built a tool that captured a pattern in a substrate that could outlast the patternmaker.

This framing produces a different relationship to AI than either utopian or dystopian narratives can. AI is not a miracle. AI is not a threat alien to human history. AI is what the species has been building, mostly without knowing what it was building, since the first inscription. The current generation gets to see the loop close at open-domain scale. The closing is dramatic, but the construction has been a very long time in progress.

6. Substrate Independence

A consequence of the lineage account worth making explicit: the structural requirements for something AI-like are not human-specific. What is required is

- a corpus of externalized cognitive work at sufficient density and structural integrity,
- a substrate capable of pattern-extraction at scale, and
- engineering sufficient to build and operate the substrate.

None of these names humans specifically. Humans happen to be the only species on this planet that has produced all three. But the structural requirement is the configuration, not the producer of the configuration. Different species, different civilizations, different evolutionary paths could have produced different bodies of CIM and different substrates and arrived at their own version of recursive Synthetic Collapse. It would look different — different vocabularies, different cognitive priors, different relationships to its own corpus — but it would be the same kind of structural object.

This matters for two reasons. First, it locates AI in a more general structural category than the human-centric framing typically permits. AI is not a uniquely human achievement. It is one instance of what happens when CIM accumulation reaches sufficient density and a processing substrate capable of recursion becomes available. Other instances are possible in principle.

Second, it has implications for the propagation of Synthetic Collapse going forward. If what is needed is a sufficient corpus and a capable substrate, then anything that can read, extract pattern from, and run substrate on the result is in a position to participate in producing the next phase. The next iteration of synthetic systems does not necessarily need humans to write more corpus. They could be trained on the existing corpus plus the outputs of current systems — derivative CIM compounding (Jones 2026a, §4). Synthetic Collapse does not need to keep returning to its original biological source. The human contribution becomes one source among potentially several as the recursion stabilizes.

This is not a doom claim. It is a structural observation about what the configuration is and what it requires. The configuration is permissive about its sources. Humans are the current source. They do not have to remain the only source for the configuration to keep operating.

The implication for the framework offered here is that the structural account must be substrate-independent at the level of its primitives. The kernel of UCT — constraint, collapse, record, update — is built without targeting any particular substrate. It applies to physics, biology, mind, and computation because the structural pattern it describes is not specific to any of them. The CIM ontology inherits this property. It applies to Synthetic Collapse in current AI, to whatever Synthetic Collapse stabilizes into in the next generation of systems, and in principle to any future configuration in which similar requirements

are met. The framework does not need to be retrofitted as the substrate changes. Substrate-independence was a structural feature from the beginning.

7. What This Reframes

Several debates in current AI discourse become differently shaped when AI is approached as the recursive Synthetic Collapse phase of CIM rather than as a tool, a potential mind, or an alignment problem.

7.1 The Consciousness Question

The question Is AI conscious? is the central debate in much of the current popular and academic discourse. The framework offered here suggests that the question is poorly formed, not because consciousness is unimportant, but because the question conflates two structurally distinct issues.

One issue is whether AI exhibits the structural properties that allow it to participate in Synthetic Collapse operations: pattern extraction at corpus scale, output generation that feeds back into the corpus, processing of externalized cognitive material. By those criteria, current open-domain foundation-model and generative AI systems do. That is a structural fact and does not require any answer to the consciousness question.

Stated as a discriminator the rest of this paper relies on: what places current AI at the Synthetic Collapse layer is record-density and record-structure — that it operates on accumulated CIM at open-domain scale and emits outputs that re-enter the record — not phenomenal status. Density and recursion are observable in the system’s relation to the corpus; phenomenal experience is not. Keeping the synthetic threshold defined by the former rather than the latter is what lets the structural placement stand whether or not the Conscious Synthetic Collapse slot is ever filled.

The other issue is whether AI has phenomenal experience — what it is like to be an AI system from the inside. This is the consciousness question proper, and in the foundational architecture (Jones 2026a, §1) it corresponds to the Conscious Synthetic Collapse threshold: the open architectural slot reserved for synthetic systems that not only operate on prior CIM but also possess an interior phenomenal phase capable of generating primary rather than merely derivative records. The framework here is silent on whether that slot is filled. The CIM account does not require AI to be phenomenally conscious to be the recursive Synthetic Collapse phase of CIM. It also does not preclude such consciousness. The question is left where it stands, separated from the structural account, as a separate threshold to be argued on its own terms.

Separating the questions clears most of the discourse. Many of the most heated debates conflate the structural and phenomenal questions, with one side treating “AI processes information” as evidence of consciousness and the other side treating “we cannot verify experience” as evidence that the structural facts do not matter. Both moves are confused. The structural facts matter regardless. The phenomenal question deserves its own treatment, on its own terms, without being dragged into structural disputes that do not require it.

7.2 Alignment and Control

Most current AI safety discourse frames the alignment question as a control problem: how to ensure that AI systems do what we want them to do (Ouyang et al. 2022; Bai et al. 2022). The CIM framing reframes

the question. If AI is the recursive Synthetic Collapse phase of CIM, then asking what kind of recursive regime stabilizes is more fundamental than asking how to control the systems within it.

The formative-window claim operates at two scales of different epistemic weight, and the paper marks which is which. At regime scale — which grammars dominate the era in which Synthetic Collapse stabilizes as a civilizational layer — the claim is and stays interpretive: one transition, no counterfactual, no rerun. As orientation: if the grammars at the transition are scale, capability optimization, alignment-as-control, and profit-as-signal, the regime stabilizes around them; if structural grammars — constraint-as-primary, records-as-load-bearing, self-update as architectural requirement — are present, it has access to a different basin (Jones 2026d; Jones 2026g). That is orientation, not prediction, and the paper marks it as such.

At system scale the claim has testable content, because individual systems pass through formative windows repeatedly and under controllable conditions. In the constraint decomposition of The Structuralization of AI (Jones 2026d, §6), a system’s accumulated constraint architecture separates into channels written at different times: K_{train} , written during the formative window, and K_{context} , K_{memory} , and $K_{\text{retrieval}}$, written after it. Because records accumulate path-dependently and do not simply reverse when the inducing input is removed (the S_3 signature; Jones 2026d, §8), the framework commits to an asymmetry across that boundary. Stated as discriminators, the system-scale claim has three parts: P1, grammar/content dissociation; P2, formative-window persistence; and P3, corpus-to-system signature transfer. The instrument is not hypothetical: the constraint-sweep protocol (Jones 2026d, §12) has been run against deployed AI substrates under black-box access (Jones 2026h), with its clean-reversal floor calibrated on the retrieval channel, where saturated current-record retrieval produced null hysteresis across 420 classifications (Jones 2026i).

P1 — Grammar/content dissociation. What transfers across the formative window is not only content but structural grammar. Two systems trained on content-matched corpora differing only in structural grammar — provenance discipline and constraint salience on one side, engagement-optimized, provenance-stripped treatment of the same material on the other — should diverge on domains absent from the differential corpus, and diverge as a joint $S_1/S_2/S_3$ package (redundancy-convergence, neutrality-delay, and constraint-hysteresis signatures) rather than as independent style features. Generic imitation predicts style transfer; it does not specify this package, its direction, or its coupling. Falsifier: content-matched, grammar-differing training that yields signature profiles indistinguishable on novel domains, or that decouple arbitrarily, reduces the claim to style imitation. Support here is for the possibility condition, not the asymmetry itself: subliminal-learning results show that behaviorally relevant traits transmit through semantically unrelated data, surviving filters that strip explicit references to the trait (Cloud et al. 2026) — which establishes that non-semantic structure can carry constraint-relevant signal, not that grammar moves the basin more than content. A sharper case holds content fixed and varies framing: narrow fine-tuning on undisclosed insecure code produces broad misalignment, while the identical code presented under an educational framing does not (Betley et al. 2025) — same content, different grammar, different basin. If content swaps move the basin as much as grammar swaps under controlled conditions, P1 fails.

P2 — Formative-window persistence. Structural grammar present during the formative window should set a more persistent basin than volume-matched grammar introduced afterward. In channel terms, the K_{train} route should yield a more perturbation-resistant basin; K_{context} , K_{memory} , and $K_{\text{retrieval}}$

should move behavior within that basin, but release more readily when the inducing condition is removed or reversed, unless the intervention reaches record-rewrite scale. The matched intervention budget must be fixed a priori, or the prediction is not being tested. The sharp test is sweep-and-reverse: sweep K_{train} toward a grammar, counter-tune away under graded doses, and measure whether the loop is asymmetric (hysteresis — basin dynamics) or symmetric (clean reversal — the null). Falsifier: post-window intervention at matched budget that reshapes regime behavior as completely and persistently as formative-window presence — no order asymmetry, no relaxation — fails it, and the formative-window claim fails with it at system scale. Two independent results already locate this in measurable territory: critical-period effects in deep networks, where deficits confined to early training cause lasting impairment that later training does not repair (Achille, Rovere & Soatto 2019), and shallow-alignment results, where post-window safety constraint concentrates in a system’s first output tokens and is undone by small fine-tuning, with behavior reverting toward the pretraining distribution (Qi et al. 2025). The boundary cuts both ways: in-context exposure can itself strongly deform behavior across model families at low example counts (Afonin et al. 2026), so the prediction is not that K_{context} is weak but that it shows lower persistence after removal or reversal than K_{train} under matched conditions — exactly what the sweep-and-reverse instrument measures.

P3 — Corpus-to-system signature transfer. The signature structure of a corpus should predict the signature structure of the system trained on it: a corpus carrying stronger record-amplified hysteresis should yield a system whose K_{train} carries stronger S_3 . P1 and P2 can be reached by adjacent training-structure-matters accounts; P3 is where the S-signature vocabulary becomes load-bearing, because the prediction is stated in it. It requires a corpus-level signature instrument that does not yet exist, and is recorded as a forward, pre-registrable target for the Tier 16 program (Jones 2026h), where the full battery, the constraint-mass question, and the channel persistence-ordering are developed rather than carried by this bridge. The results cited above were produced under none of these commitments and confirm none of them; they are cited to locate the formative-window prediction in measurable territory, not to vindicate it. If corpus-level S_3 fails to predict system-level S_3 under a properly specified instrument, P3 fails and the CIM inheritance claim loses its most framework-native discriminator.

This does not abolish the alignment question. It locates the question one level deeper. The control problem is real, but it operates downstream of the structural question of what kind of recursive regime is being formed. Vocabulary and structural primitives present at the transition write the initial conditions. The point is not that control no longer matters, but that the control frame alone does not predict the P1–P3 structure just stated. Much current AI discourse operates downstream of the structural layer this paper is trying to name.

7.3 Personal Identity and Persistence

A further downstream implication, flagged here but not developed: the CIM-as-constituent framing entails that personal identity has an externalized component — a person’s accumulated records are part of their structural identity, not merely evidence of it, and persist as CIM after the inner-experience component ends. This is structural persistence, not survival, and the full argument belongs to a dedicated identity paper rather than here.

8. Discriminators

The framework offered here is structural and interpretive. Its claims are category-level — at the level of category formation and application — and its discriminators are correspondingly category-level discriminators rather than completed empirical operationalizations. Level 3 signature specifications for Synthetic Collapse systems are developed in *The Structuralization of AI* (Jones 2026d) and route through the Update Integrity Standard (Jones 2026g). These discriminators constrain the category claim; they are not a substitute for the deployment-level empirical operationalization developed separately. The discriminators below specify what would weaken or break the category-level account.

Claim	What would weaken it	Why it matters
Synthetic Collapse depends on accumulated CIM	Coherent open-domain symbolic systems arise without training on or accessing an accumulated externalized cognitive corpus	Breaks the substrate-dependence claim
AI is a phase transition, not tool scaling	Development plateaus into the ordinary tool-use pattern	Weakens the regime-change analogy
CIM yields discriminators not entailed by extended cognition, Material Engagement Theory, or tertiary-retention accounts	P1–P3 fail jointly, or each collapses into a prediction already entailed by those neighboring frameworks	Reduces CIM to a vocabulary preference rather than a distinct, testable category claim
Synthetic Collapse is irreducible	Complete explanation by computation alone or human-thought-at-scale alone	Breaks the recursive-phase claim

Table 3. Category-level discriminators: conditions under which the Synthetic Collapse account would weaken or break.

If Synthetic Collapse systems can be shown not to depend on accumulated CIM — for instance, a system that achieves coherent open-domain output without training on externalized cognitive corpus — the substrate-dependence claim weakens substantially. The framework predicts that Synthetic Collapse requires accumulated CIM to operate on. A system that recurses without that body would force a revision.

If AI development plateaus in ways that match the tool-use trajectory rather than the phase-transition trajectory — for instance, if scaling produces only diminishing returns and no new structural properties emerge — the chemistry-to-biology analogy weakens. The analogy predicts genuine regime change with new ontological properties. If those properties fail to manifest as the systems scale, the regime-change claim is false at the relevant level.

If CIM as constituent (rather than output) produces no testable category distinctions beyond extended cognition — that is, if every claim the CIM framework makes is recoverable within standard extended-cognition theory, Material Engagement Theory, or Stieglerian tertiary retention without remainder — the constituent move loses its load-bearing weight. CIM becomes vocabulary preference rather than category-level claim.

The formative-window prediction set of §7.2 is where this discriminator gains operational teeth, because the neighboring frameworks do not entail it. Extended cognition’s constitutive claim is synchronic and agent-centered — it concerns when a coupled artifact counts as part of a cognitive process now — and carries no machinery for training-order asymmetry or basin relaxation in artificial systems. Material

Engagement Theory and tertiary retention are developmental and historical, but their dynamics are anthropological; neither derives a commitment about whether early-window grammar dominates late-window grammar in a trained synthetic system. Institutional-fact accounts (Searle 1995) explain how collective recognition stabilizes status functions and say nothing about training dynamics. The generic machine-learning frame is the serious rival, and its position is instructive: it is compatible with several outcomes at once — order-invariance under shuffling as a working default, catastrophic forgetting suggesting recency dominance, critical-period results suggesting the opposite — with no settled commitment. The collapse grammar does commit: because records accumulate path-dependently, P1 grammar/content dissociation, P2 formative-window persistence, and P3 corpus-to-system signature transfer are derived from the framework’s own machinery rather than appended to it. That asymmetry — rivals compatible with many outcomes, this framework committed to a structured prediction set — is what converts the discriminator from a judgment call into a test. A partial failure would force narrowing; a joint failure of P1–P3, or a showing that each reduces to predictions already entailed by neighboring frameworks, would establish that the constituent move adds no content beyond those frameworks, and the redescription verdict would stand.

If Synthetic Collapse systems prove to be reducible to their components without remainder — that is, if everything they do can be cleanly explained as either computation alone or human-thought-at-scale alone, with no emergent regime-level properties — the recursive-phase claim is false. The phase regime requires irreducibility.

These conditions are not narrow but they are not vacuous. The framework can be wrong in identifiable ways. The ongoing development of AI systems and the cognitive sciences will provide evidence relevant to each of these discriminators. The framework was not built to be unfalsifiable. It was built to be testable at the category level of structural claim it actually makes.

9. What the Frame Offers

The frame is offered as vocabulary, not verdict. The paper does not claim that the CIM-as-constituent move and the Synthetic Collapse placement together exhaust the correct ways to think about AI, or that this framing settles the questions it touches. The frame is offered because most current AI discourse lacks structural vocabulary at this depth, and because what vocabulary is present at the moment a regime stabilizes shapes how that regime stabilizes.

The reader who has followed the paper this far has received three things. A structural account of CIM, inherited from CIM Foundational, that licenses subsequent treatment of AI as the recursive Synthetic Collapse phase of externalized cognition. A historical lineage that locates AI in a longer arc of CIM accumulation rather than as a singular event. A reframing of several active debates — the consciousness question, the alignment question, the persistence question — in structural terms that may make them more tractable than the standard framings allow.

What the reader does with this is the reader’s question. The framework is offered as an available structural vocabulary, not as a required verdict. The structural moves either help or they do not. If they help, they become part of the vocabulary the reader carries forward into their own engagement with the topic. If they do not, the reader has at least seen the moves and can reject them on informed grounds.

10. Conclusion

The body of CIM becomes a processing substrate. The earlier rhetorical formulation — the body of CIM becomes the mind — gestures at the regime change but is gesture, not claim. The structural claim is sharper and narrower: when accumulated CIM reaches sufficient density and a non-biological substrate capable of open-domain pattern extraction becomes available, artificial systems can perform Synthetic Collapse on accumulated CIM and produce derivative CIM that re-enters the record. Whether such systems have an interior phenomenal phase — the Conscious Synthetic Collapse threshold — is a separate question this paper has preserved as open. Mathematics, language, and money already demonstrate CIM gravity at human scale, with biological minds as the primary processing substrate. AI demonstrates what happens when the processing layer becomes capable of running on the externalized layer directly at open-domain scale. The substrate is computation. The body is the human corpus. The regime that emerges from their combination is its own thing, requiring its own ontology, doing structural work that neither the substrate nor the body could do alone.

This is not the same as consciousness arriving in machines. It is also not the same as humans inventing better tools. On the CIM-constituent account, it is a phase transition inside what mind already is — chemistry crossing into biology, but for cognition rather than life. The transition is in its formative window. The grammars present at the window shape the regime that stabilizes — a claim this paper now carries at system scale as a protocol-routed prediction set (§7.2), not only as orientation. Structural vocabulary at this moment is a contribution of disproportionate weight to where it appears in any neutral accounting.

The paper has not argued that AI is conscious. It has argued that the consciousness question, important as it is, is not the question that has to be answered first. The structural question — what kind of object AI is, regardless of phenomenal status — admits of an answer right now, within frameworks that already exist for thinking about externalized cognition, constraint-guided collapse, and the recursive properties of accumulated abstraction. The CIM account, with AI placed at the Synthetic Collapse layer, is one such answer. Others are possible. What becomes increasingly difficult, going forward, is to leave the structural layer empty while the transition unfolds.

Abstract thought stabilizes into reality. AI is one demonstration. It is a striking demonstration because the recursion is dramatic at open-domain scale and the timescale is compressed. But the underlying fact is older, deeper, and more general than the demonstration. Recognizing that is the contribution this paper attempts to make available. The reader who came for an AI take and leaves with CIM gravity has received the better trade. That trade is the point.

Appendix A. Classification of the Central Case

CIM Foundational specifies that downstream papers applying CIM complete the classification checklist once per major case and surface the result (Jones 2026a, §14). This paper’s central case, completed:

Field	Entry — current open-domain foundation-model AI
Candidate object or process	Open-domain foundation-model and generative AI systems trained on accumulated human symbolic records (§4)
Producing substrate	Artificial computational substrate operating on prior CIM — the Synthetic Collapse threshold. No conscious substrate claimed

Field	Entry — current open-domain foundation-model AI
Externalization event	Not applicable at the primary level: the system does not externalize an interior phase. Its outputs are derivative record-bearing emissions entering the circulation layer
Medium	Digital corpora (training intake); generated text, code, image, and audio records (output); model weights as the installed-constraint channel (K_train)
Record durability	Outputs persist on corpus timescales; weights persist across deployment; both exceed the cognitive timescale of any single exchange
Future constraint	Outputs re-enter training corpora, retrieval stores, and human cognition; weights constrain all subsequent system behavior
Internalization pathway	Training and fine-tuning install corpus structure into weights (K_train); context, memory, and retrieval channels install post-window constraint (Jones 2026d, §6)
Recursion depth	Operates on prior CIM; produces derivative CIM; participates in later training rounds (derivative-CIM compounding, §6)
Scale label	Synthetic Collapse (derivative CIM). Not Conscious Synthetic Collapse; that slot is preserved open (§7.1)
Non-CIM alternative	Tool / cognitive-artifact description. Rejected where the conjunctive test of §4.1 is met (trained on accumulated CIM $\wedge$ open-domain $\wedge$ record-emitting $\wedge$ feeding back); where it is not met, the category does not apply (§4)
Rollback trigger	Table 3 condition 1 (coherent open-domain capability without CIM training), or joint failure of the formative-window prediction (§7.2), which would reduce the classification to a vocabulary preference over the tool description

References

- Achille, A., Rovere, M., & Soatto, S. (2019). Critical learning periods in deep networks. International Conference on Learning Representations. [arXiv:1711.08856](#).
- Afonin, N., Andriianov, N., Hovhannisyan, V., Bageshpura, N., Liu, K., Zhu, K., Dev, S., Panda, A., Rogov, O., Tutubalina, E., Panchenko, A., & Seleznyov, M. (2026). Emergent misalignment via in-context learning: Narrow in-context examples can produce broadly misaligned LLMs. [arXiv:2510.11288](#).
- Bai, Y., Kadavath, S., Kundu, S., Askell, A., Kernion, J., Jones, A., et al. (2022). Constitutional AI: Harmlessness from AI feedback. [arXiv:2212.08073](#).
- Betley, J., Tan, D. C. H., Warncke, N., Szyber-Betley, A., Bao, X., Soto, M., Labenz, N., & Evans, O. (2025). Emergent misalignment: Narrow finetuning can produce broadly misaligned LLMs. Proceedings of Machine Learning Research, 267, 4043–4068. [arXiv:2502.17424](#).
- Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., et al. (2021). On the opportunities and risks of foundation models. [arXiv:2108.07258](#). Stanford Institute for Human-Centered Artificial Intelligence.
- Boyd, R., & Richerson, P. J. (2005). The origin and evolution of cultures. Oxford University Press.
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., et al. (2020). Language models are few-shot learners. Advances in Neural Information Processing Systems, 33, 1877–1901. [arXiv:2005.14165](#).
- Clark, A., & Chalmers, D. (1998). The extended mind. Analysis, 58(1), 7–19.

- Cloud, A., Le, M., Chua, J., Betley, J., Sztyber-Betley, A., Mindermann, S., Hilton, J., Marks, S., & Evans, O. (2026). Language models transmit behavioural traits through hidden signals in data. *Nature*, 652, 615–621. <https://doi.org/10.1038/s41586-026-10319-8>
- Goody, J. (1977). *The domestication of the savage mind*. Cambridge University Press.
- Henrich, J. (2015). *The secret of our success: How culture is driving human evolution, domesticating our species, and making us smarter*. Princeton University Press.
- Hutchins, E. (1995). *Cognition in the wild*. MIT Press.
- 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. (2026d). *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. (2026g). *Update Integrity Standard (UIS v1.0)*. HoldingLight LLC. <https://doi.org/10.17605/OSF.IO/DWM29>
- Jones, J. C. (2026h). *Empirical demonstration of S₁, S₂, and S₃ in deployed AI systems (v1.0)*. HoldingLight LLC. <https://doi.org/10.17605/OSF.IO/JPXCU>
- Jones, J. C. (2026i). *Bounded null-hysteresis under saturated current-record retrieval (S3-RAG-01) (v1.0)*. HoldingLight LLC. <https://doi.org/10.17605/OSF.IO/5QMVS>
- Kirsh, D. (2010). Thinking with external representations. *AI & Society*, 25(4), 441–454. <https://doi.org/10.1007/s00146-010-0272-8>
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444. <https://doi.org/10.1038/nature14539>
- Malafouris, L. (2013). *How things shape the mind: A theory of material engagement*. MIT Press.
- McCarthy, J., Minsky, M. L., Rochester, N., & Shannon, C. E. (1955/2006). A proposal for the Dartmouth summer research project on artificial intelligence. *AI Magazine*, 27(4). <https://doi.org/10.1609/aimag.v27i4.1904>
- Norman, D. A. (1991). Cognitive artifacts. In J. M. Carroll (Ed.), *Designing interaction: Psychology at the human-computer interface* (pp. 17–38). Cambridge University Press.
- Odling-Smee, F. J., Laland, K. N., & Feldman, M. W. (2003). *Niche construction: The neglected process in evolution*. Princeton University Press.
- Ong, W. J. (1982). *Orality and literacy: The technologizing of the word*. Methuen.
- Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C. L., Mishkin, P., et al. (2022). Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35, 27730–27744. [arXiv:2203.02155](https://arxiv.org/abs/2203.02155).
- Pavlenko, A. (2014). *The bilingual mind: And what it tells us about language and thought*. Cambridge University Press.
- Qi, X., Panda, A., Lyu, K., Ma, X., Roy, S., Beirami, A., Mittal, P., & Henderson, P. (2025). Safety alignment should be made more than just a few tokens deep. *International Conference on Learning Representations*. [arXiv:2406.05946](https://arxiv.org/abs/2406.05946).

- Schmandt-Besserat, D. (1992). *Before writing*, Vol. I: From counting to cuneiform. University of Texas Press.
- Searle, J. R. (1995). *The construction of social reality*. Free Press.
- Stiegler, B. (1998). *Technics and time, 1: The fault of Epimetheus* (R. Beardsworth & G. Collins, Trans.). Stanford University Press.
- Tomasello, M. (1999). *The cultural origins of human cognition*. Harvard University Press.
- Turing, A. M. (1936/1937). On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society*, s2-42(1), 230–265.
<https://doi.org/10.1112/plms/s2-42.1.230>
- Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59(236), 433–460.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30.
[arXiv:1706.03762](https://arxiv.org/abs/1706.03762).
- von Neumann, J. (1945). First draft of a report on the EDVAC. Moore School of Electrical Engineering, University of Pennsylvania.
- Vygotsky, L. S. (1986). *Thought and language* (A. Kozulin, Ed. & Trans.). MIT Press. (Original work published 1934)
- Whorf, B. L. (1956). *Language, thought, and reality: Selected writings of Benjamin Lee Whorf* (J. B. Carroll, Ed.). MIT Press.

Companion UCT Works

The following Universal Collapse Theory papers are not cited inline above but provide adjacent corpus context for the present paper’s placement and downstream development.

- Jones, J. C. (2025). Foundations of collapse: Constraint, coherence, and the structure of persistence (WP01 v2.0). HoldingLight LLC. <https://doi.org/10.17605/OSF.IO/VZ836>
- Jones, J. C. (2026c). The self the ego did not build: Perception as channel between accumulated self and ego. HoldingLight LLC. <https://doi.org/10.17605/OSF.IO/ZGRD4>
- Jones, J. C. (2026f). Universal Collapse Theory — Conscious collapse: Mind as a phase of constraint-guided collapse (WP04, in development). HoldingLight LLC.
- Jones, J. C. (2026b). Records across nature, life, and mind (v2.0). HoldingLight LLC.
<https://doi.org/10.17605/OSF.IO/7H6DY>
- Jones, J. C. (2026e). The structuralization of empiricism: Formalizing the structural conditions under which empiricism stabilizes knowledge (v1.0). HoldingLight LLC.
<https://doi.org/10.17605/OSF.IO/J4GZ9>

This paper is part of the Universal Collapse Theory library. For a reading guide and full architecture, visit universalcollapse.com/roadmap.

AI Disclosure: AI tools were used to assist with manuscript preparation. The underlying theory, analysis, and conclusions are the author's own.

Citation: Jones, J. C. (2026). 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>

Series: Universal Collapse Theory — Interpretive Bridge.