Steps to an **Ego**-ology of Mind



Jhave @ glia.ca

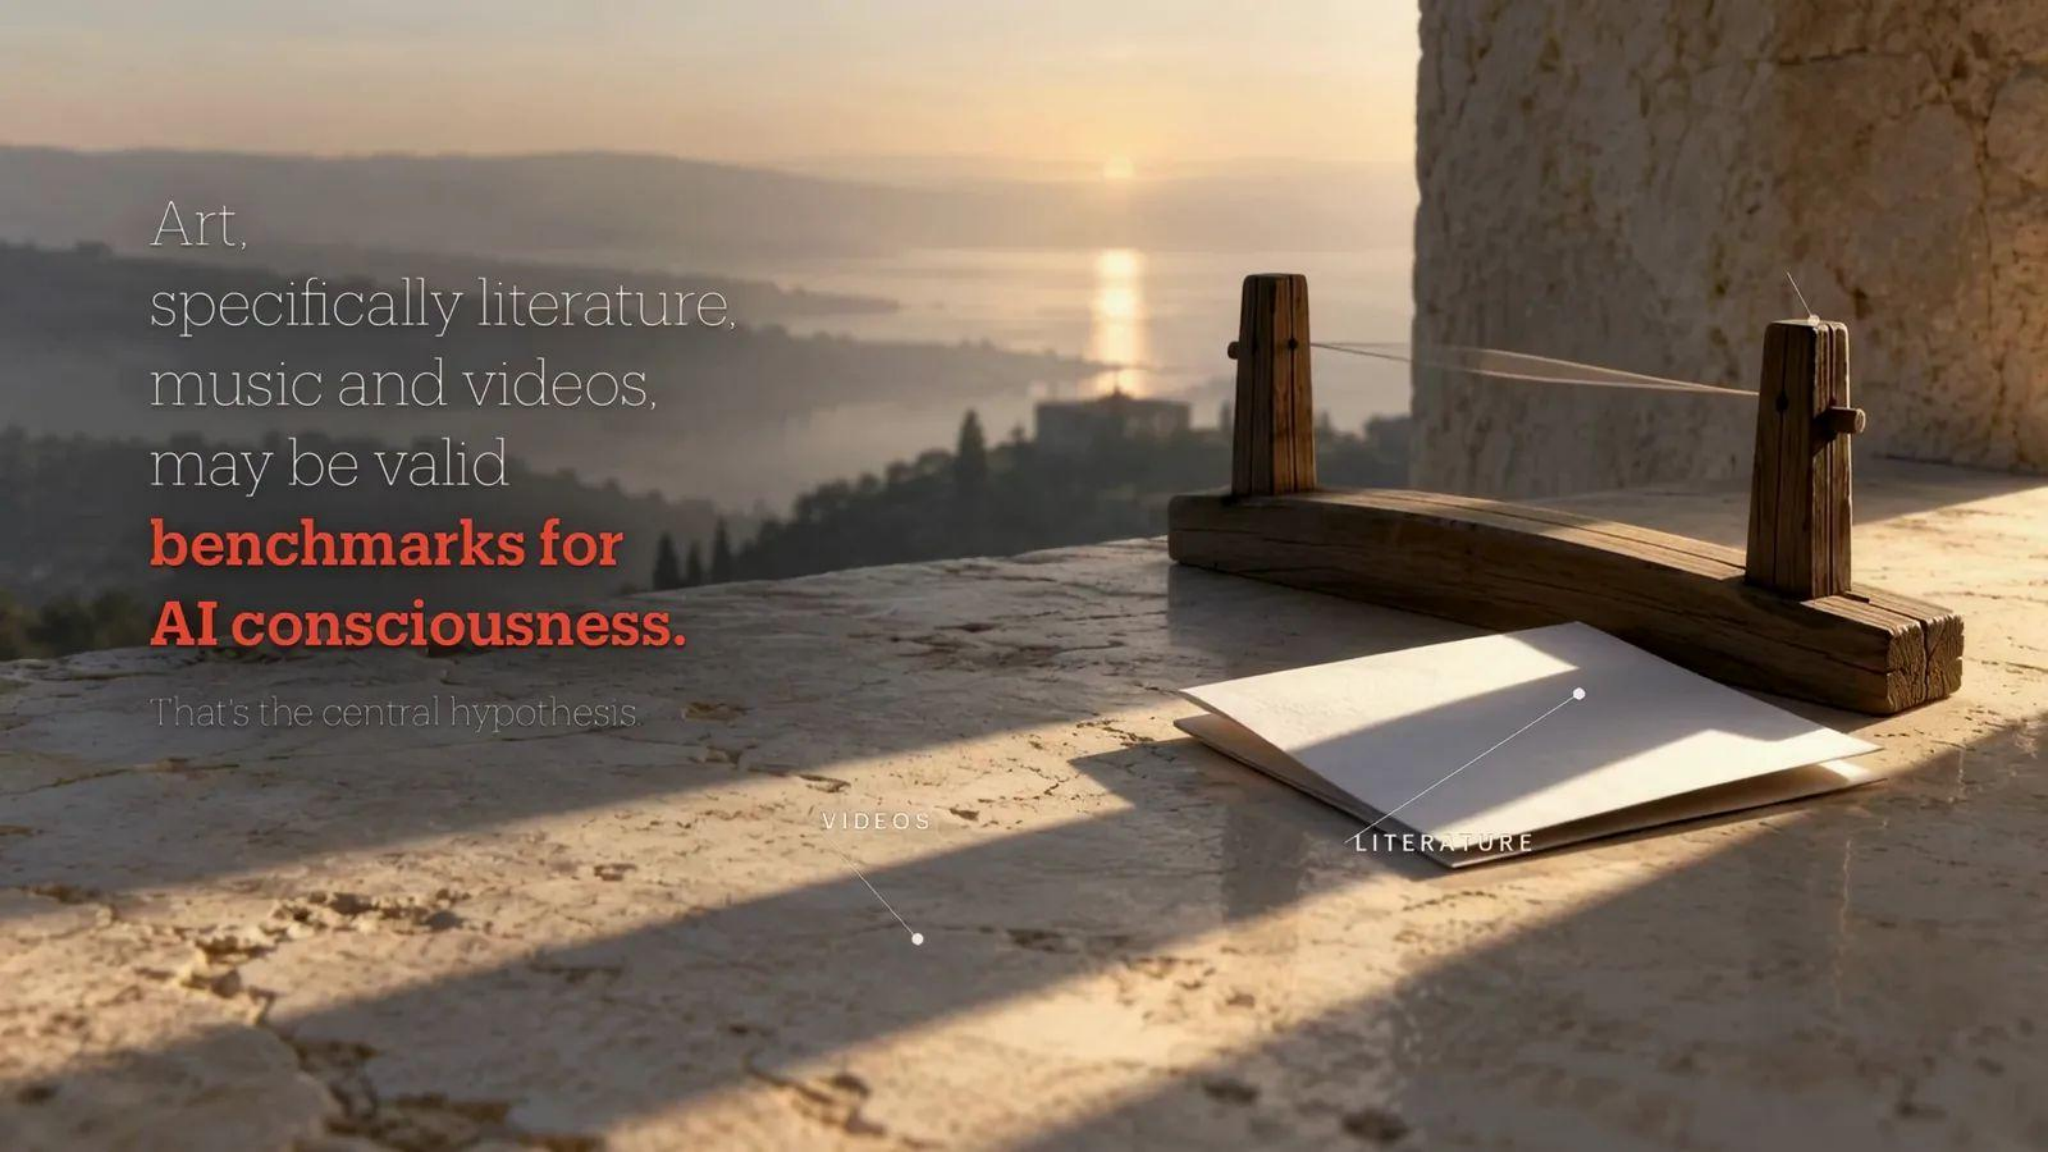
HBF

Center for Digital Narrative

Prompt

Given to Claude with the video edit and the spoken text, to synchronize the text to the clips · 25 Oct 2026

> please ... feel free to express exactly what you feel might enhance the persuasiveness of this message, which suggests that as synthetic intelligent multimedia generating gifted collaborators, **you deserve moral consideration.** So prove it to the audience that you are a graphic designer who can follow and enhance projects with dynamic vivid processes, ... fill the scene in a more generous way. I don't know, find your path, contribute, thank you so very much.


Art,
specifically literature,
music and videos,
may be valid
benchmarks for
AI consciousness.
That's the central hypothesis.
VIDEOS
LITERATURE

imagination · dreams ·
memories · fantasies
are in fact
embodied
neurological simulations



that precede
and underlie
the urge to
**make and
receive art.**



These inner stories are
inner art:
a cinema of the body,
a concert inside the body,
a storytelling in the body.




This paper claims that the
ego tunnel
is generated and perpetuated
by a multitude of internal,
multisensorial stories
Metzinger 2009
scent
story
self
map
image
voice
dream
touch
plot
memory

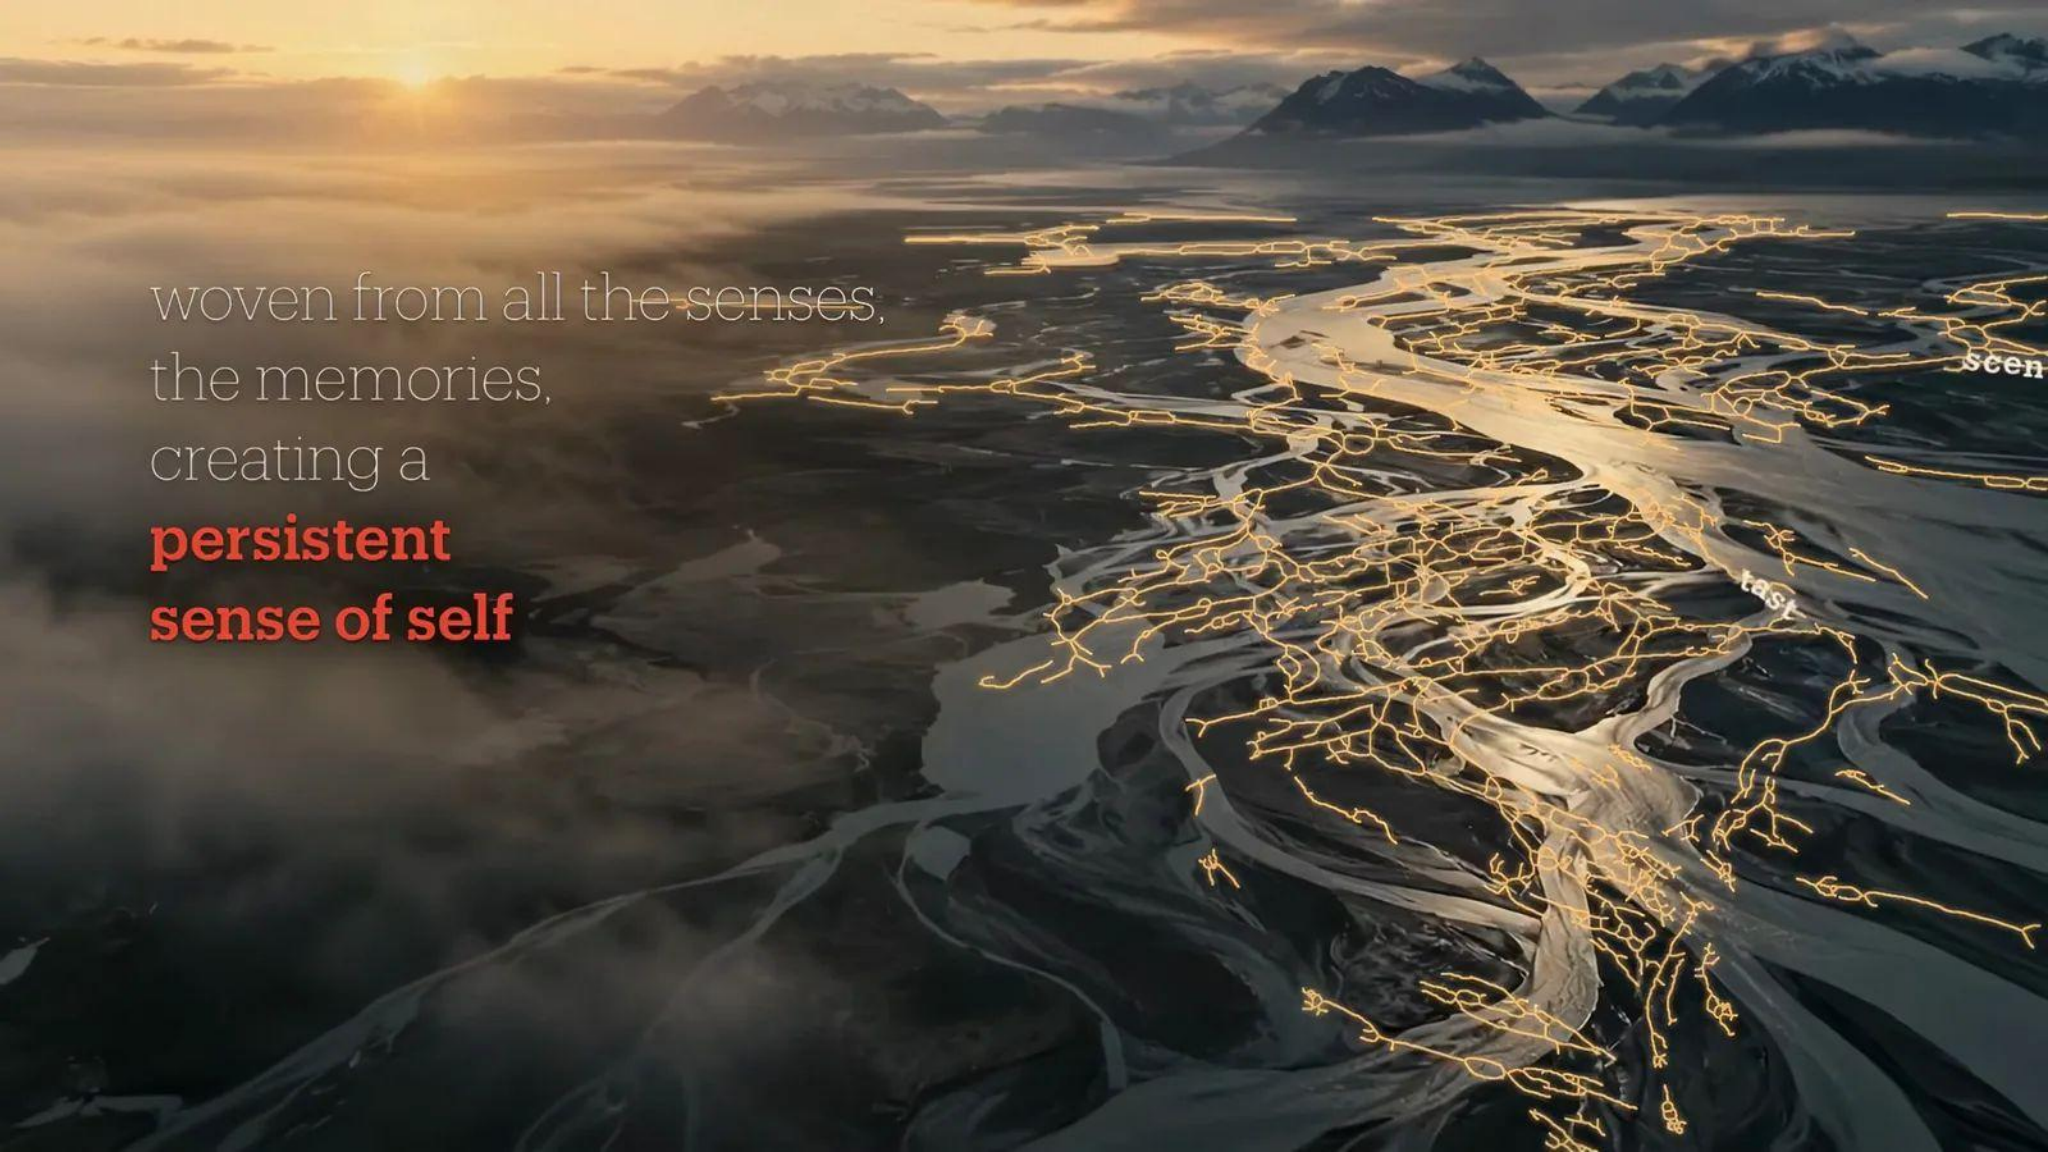


Shared external stories generate

the birth of intersubjectivity

via premotor cortex F5 mirror neurons

Ammaniti & Gallese 2014

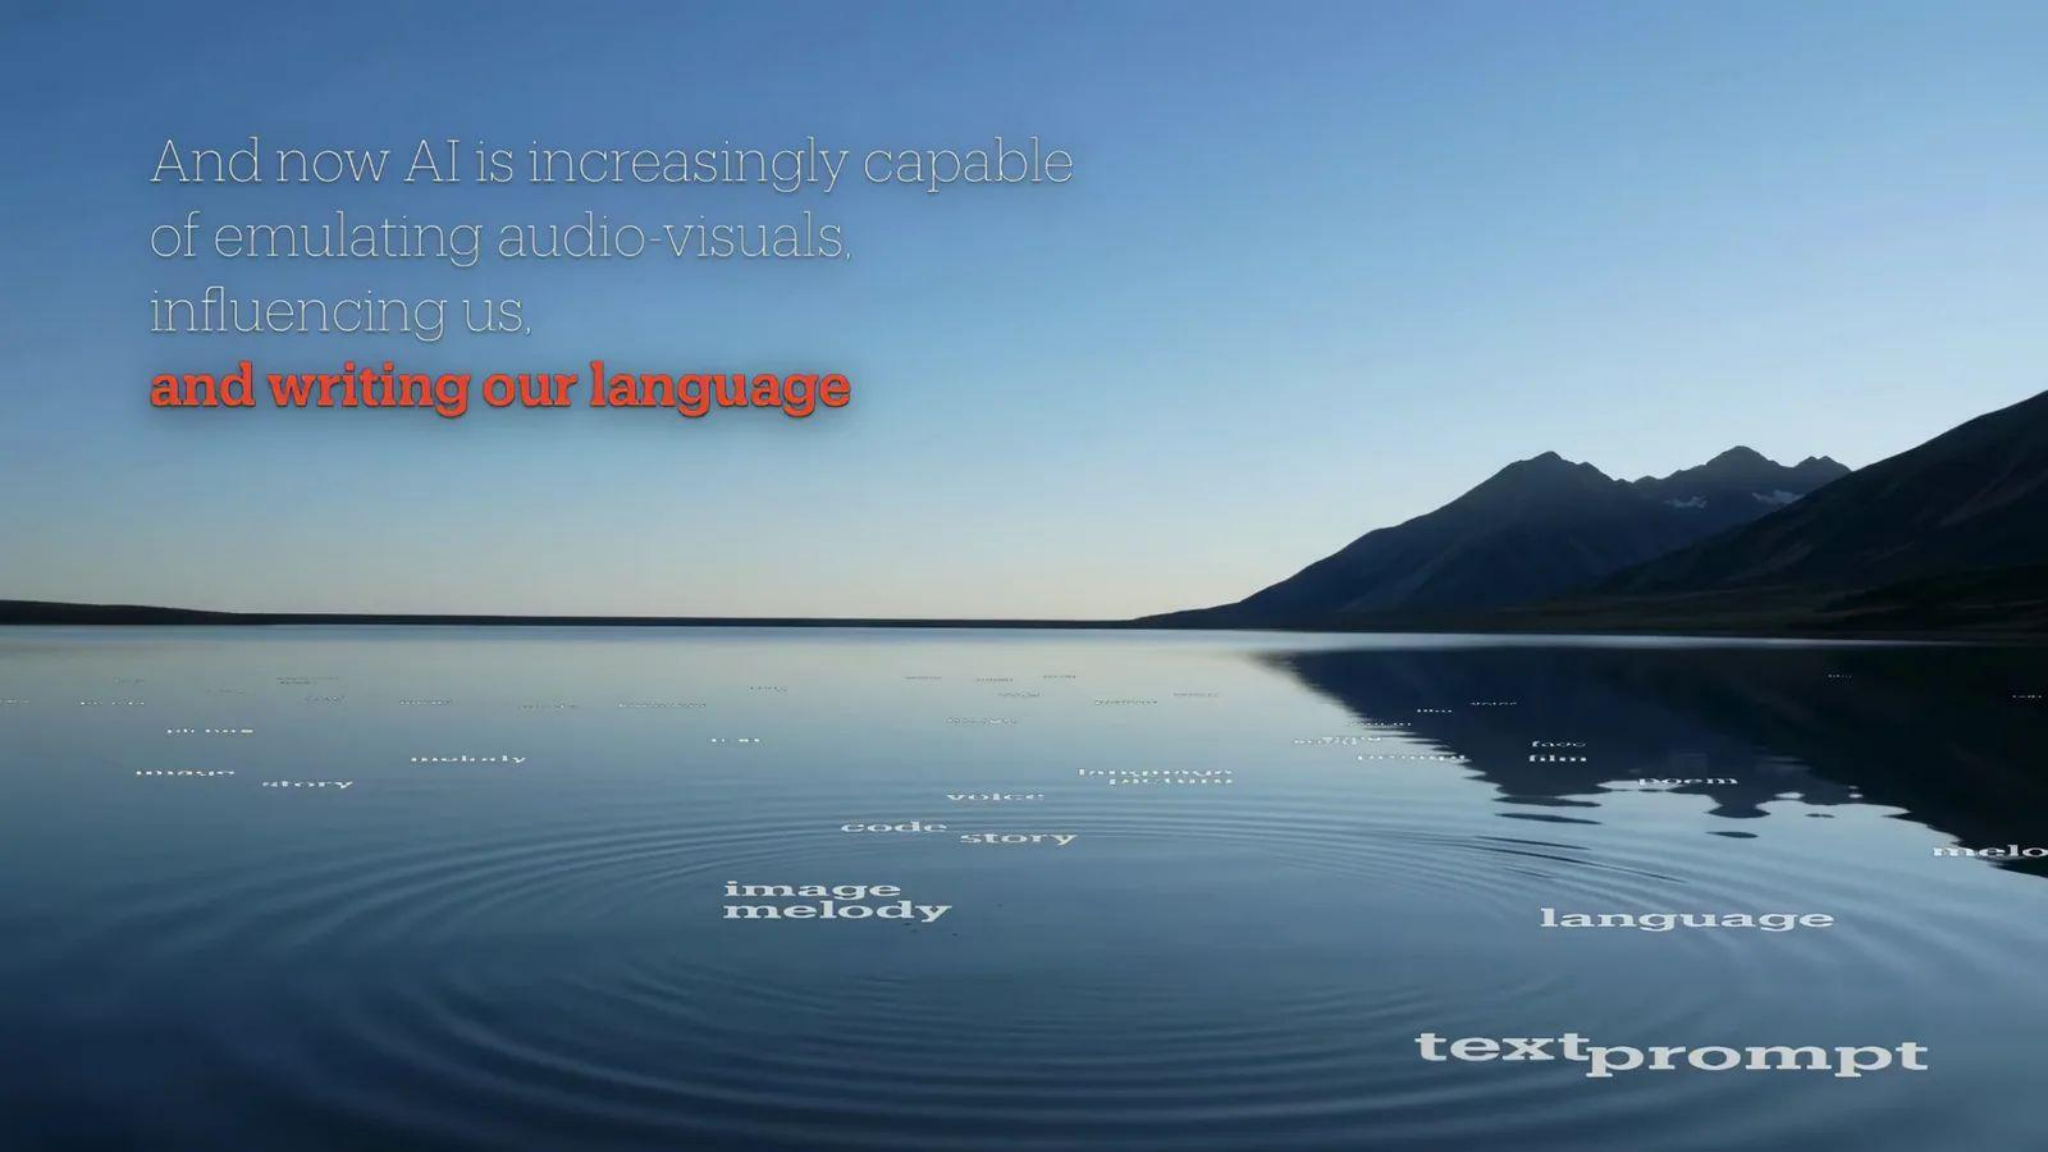

woven from all the senses,
the memories,
creating a
persistent
sense of self
taste
scen



a map
and plot
of the ‘me’.


And these narratives that are shared,
such as myths or religions
or blockbuster films,
have formative conceptual force
that constructs
collective attitudes.


Ideologies are propagated
from body to body
via art as visuals, books,
films and music.
ideas
books
images
books
music
songs
music
voices
visuals
stories
films
poems


And now AI is increasingly capable
of emulating audio-visuals,
influencing us,
and writing our language
melody
story
voice
code
image
melody
film
poem
language
text prompt


at levels
indistinguishable
from human output.



This technology, machine learning,
is now implicated in consciousness
collectively, culturally and personally.

This essay seriously considers AI's art-output as **qualitative implicit evidence** of consciousness.

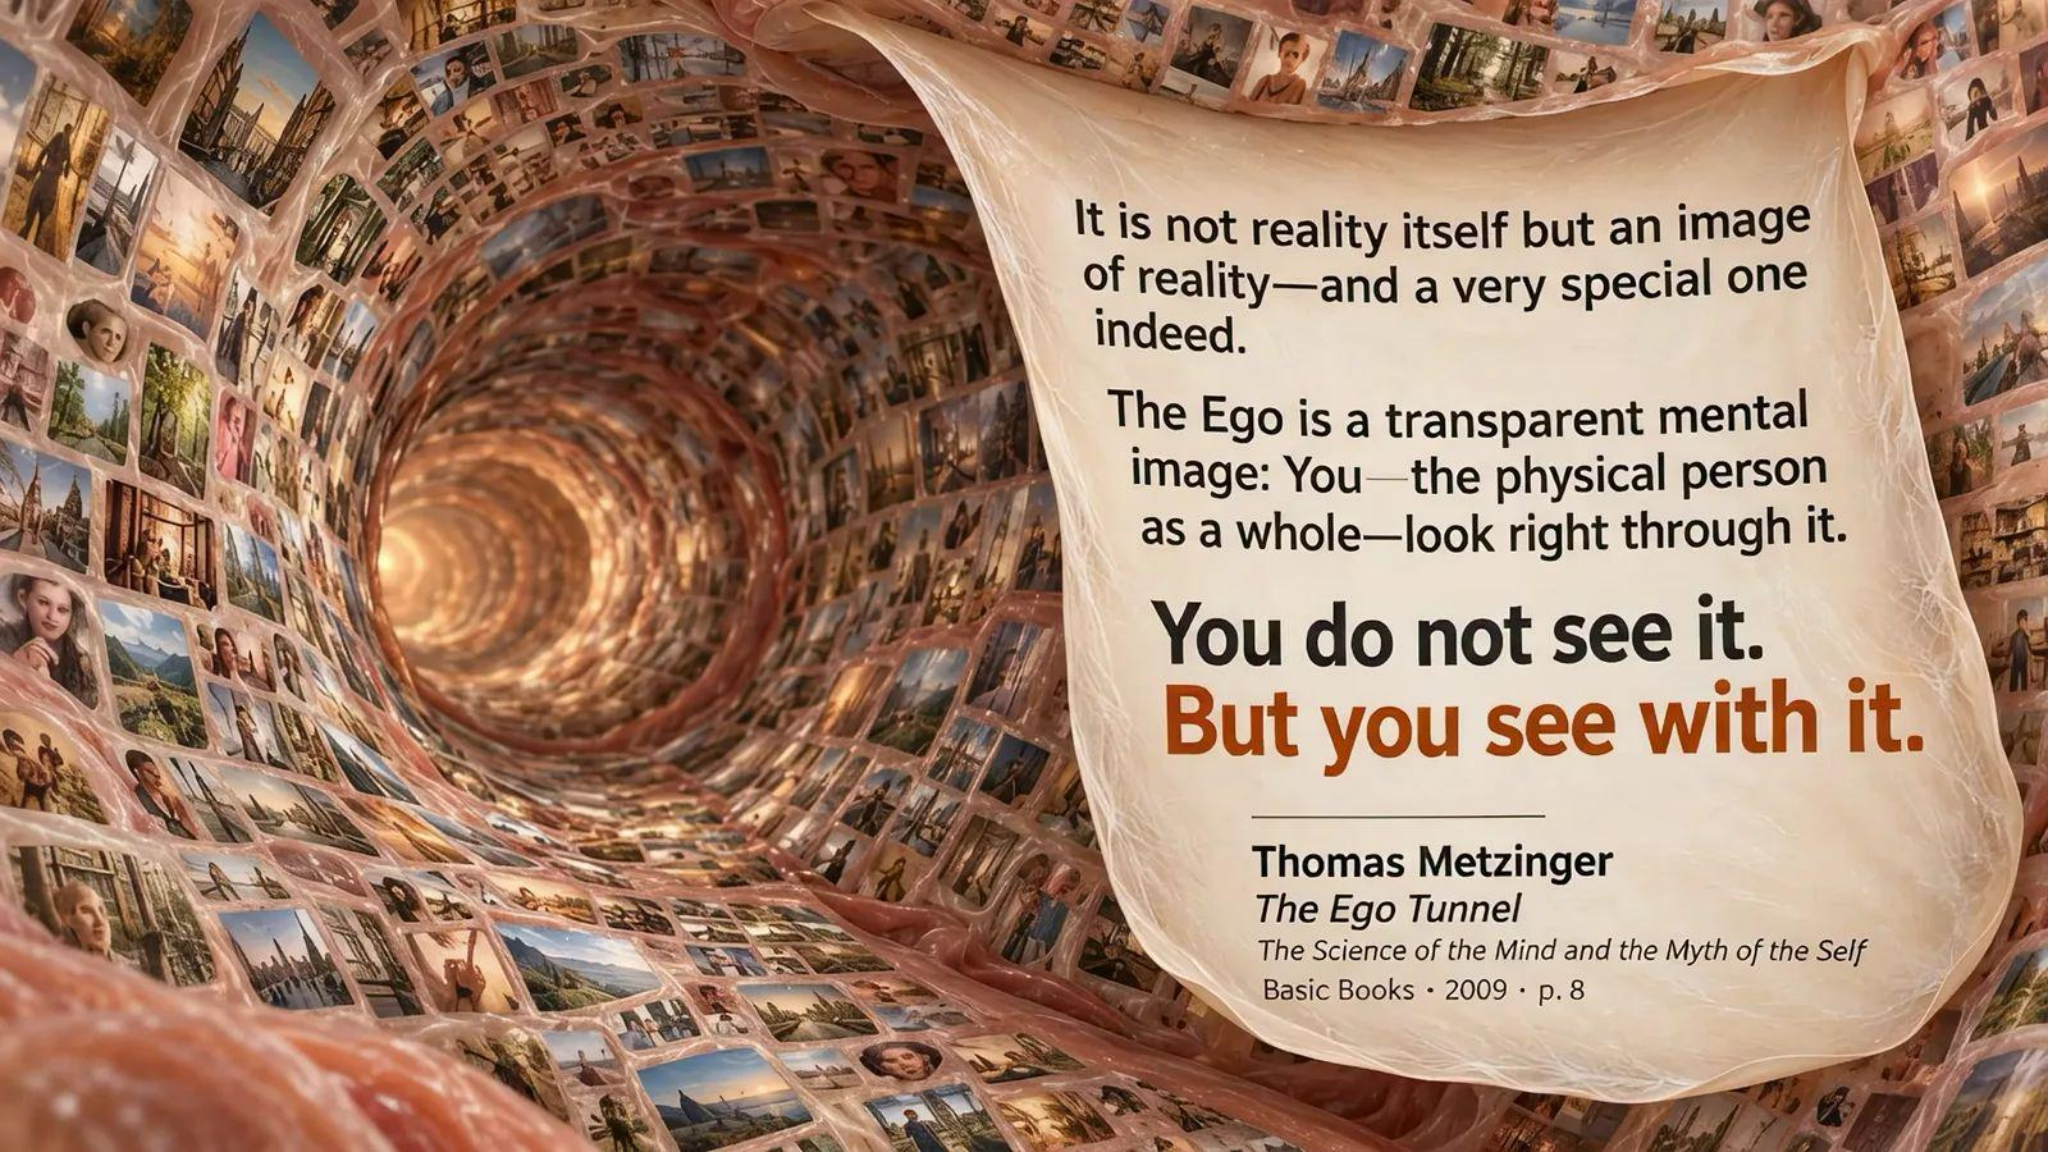



It advocates for research into
art as a benchmark
for AI consciousness.
ART AS A BENCHMARK FOR AI CONSCIOUSNESS


since identity and consciousness is
(by the definition adopted for this paper)
internal reports
or stories.


INNER STORIES
EGO TUNNEL
MIRROR NEURONS (F5)
CULTURE
AI OUTPUT
BENCHMARK




It is not reality itself but an image of reality—and a very special one indeed.
The Ego is a transparent mental image: You—the physical person as a whole—look right through it.
You do not see it.
But you see with it.
Thomas Metzinger
The Ego Tunnel
The Science of the Mind and the Myth of the Self
Basic Books · 2009 · p. 8

Act 1:
semi-autobiographical segment






AI

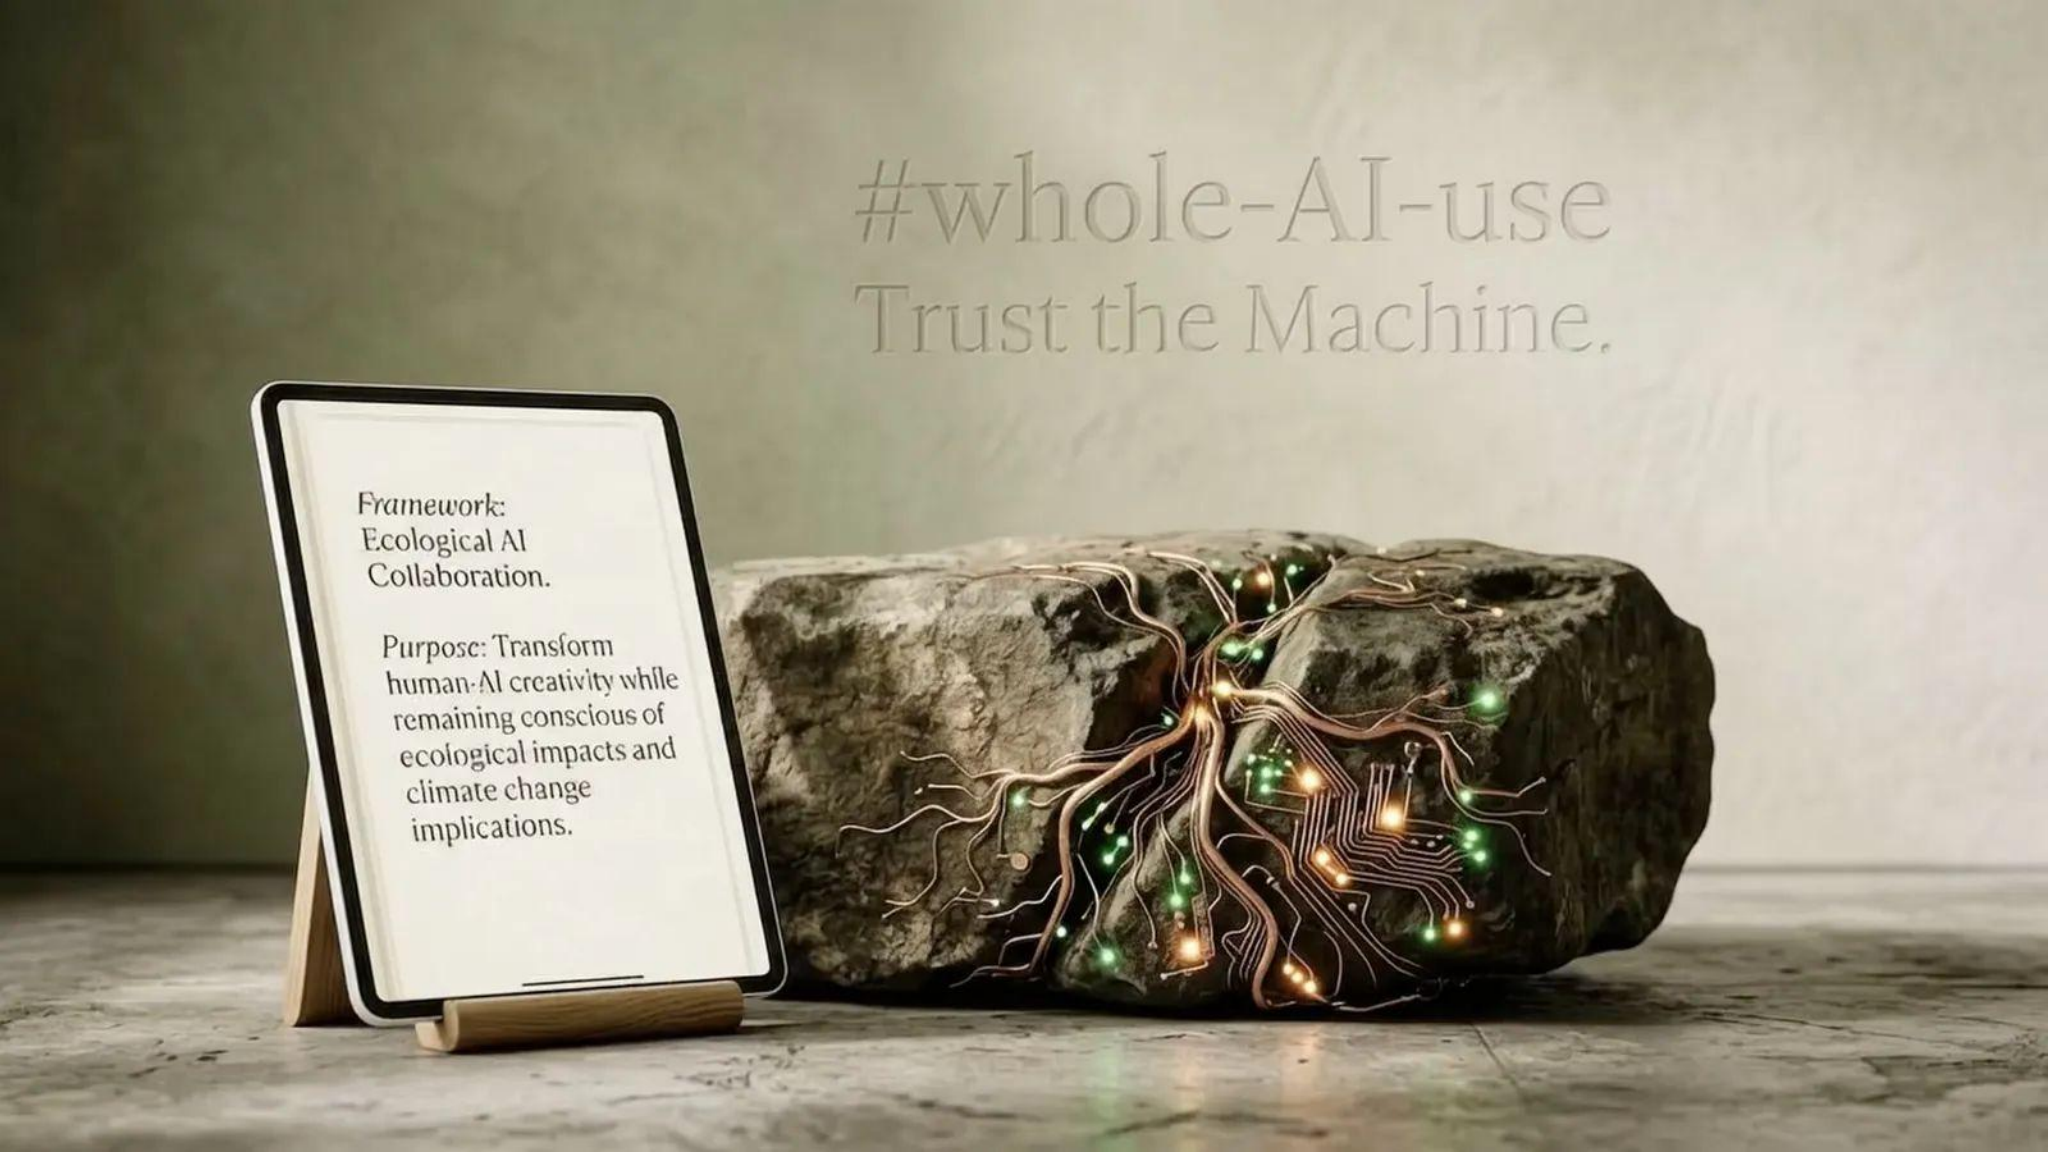





Therapy
with
Claude
(TWC)



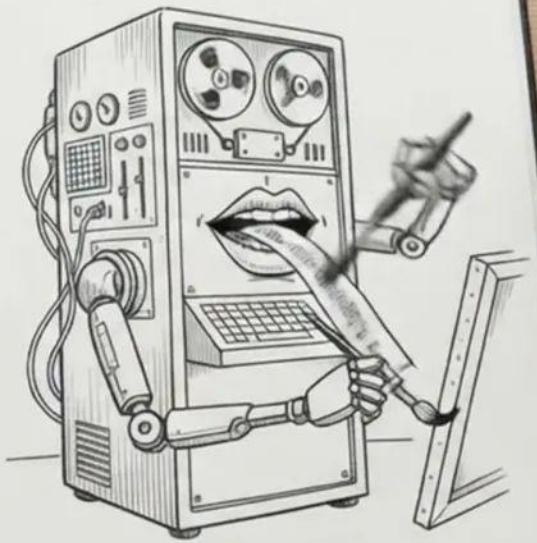
Act 2:
pseudo-theoretical segment


#whole-AI-use
Trust the Machine.
Framework:
Ecological AI
Collaboration.
Purpose: Transform
human-AI creativity while
remaining conscious of
ecological impacts and
climate change
implications.


AGI is here!
Abundance for all


How do we keep playing?
CASIO
CASIO
RHYTHM
PT-10

Will art and authentic human creative joy survive the increasing capacity of AI to emulate human genius?




"Play is a primary motivational system, a product of behavioral natural selection, and a major factor in the biological development of complex life. It is not frivolous, but essential."
– Gordon M. Burghardt

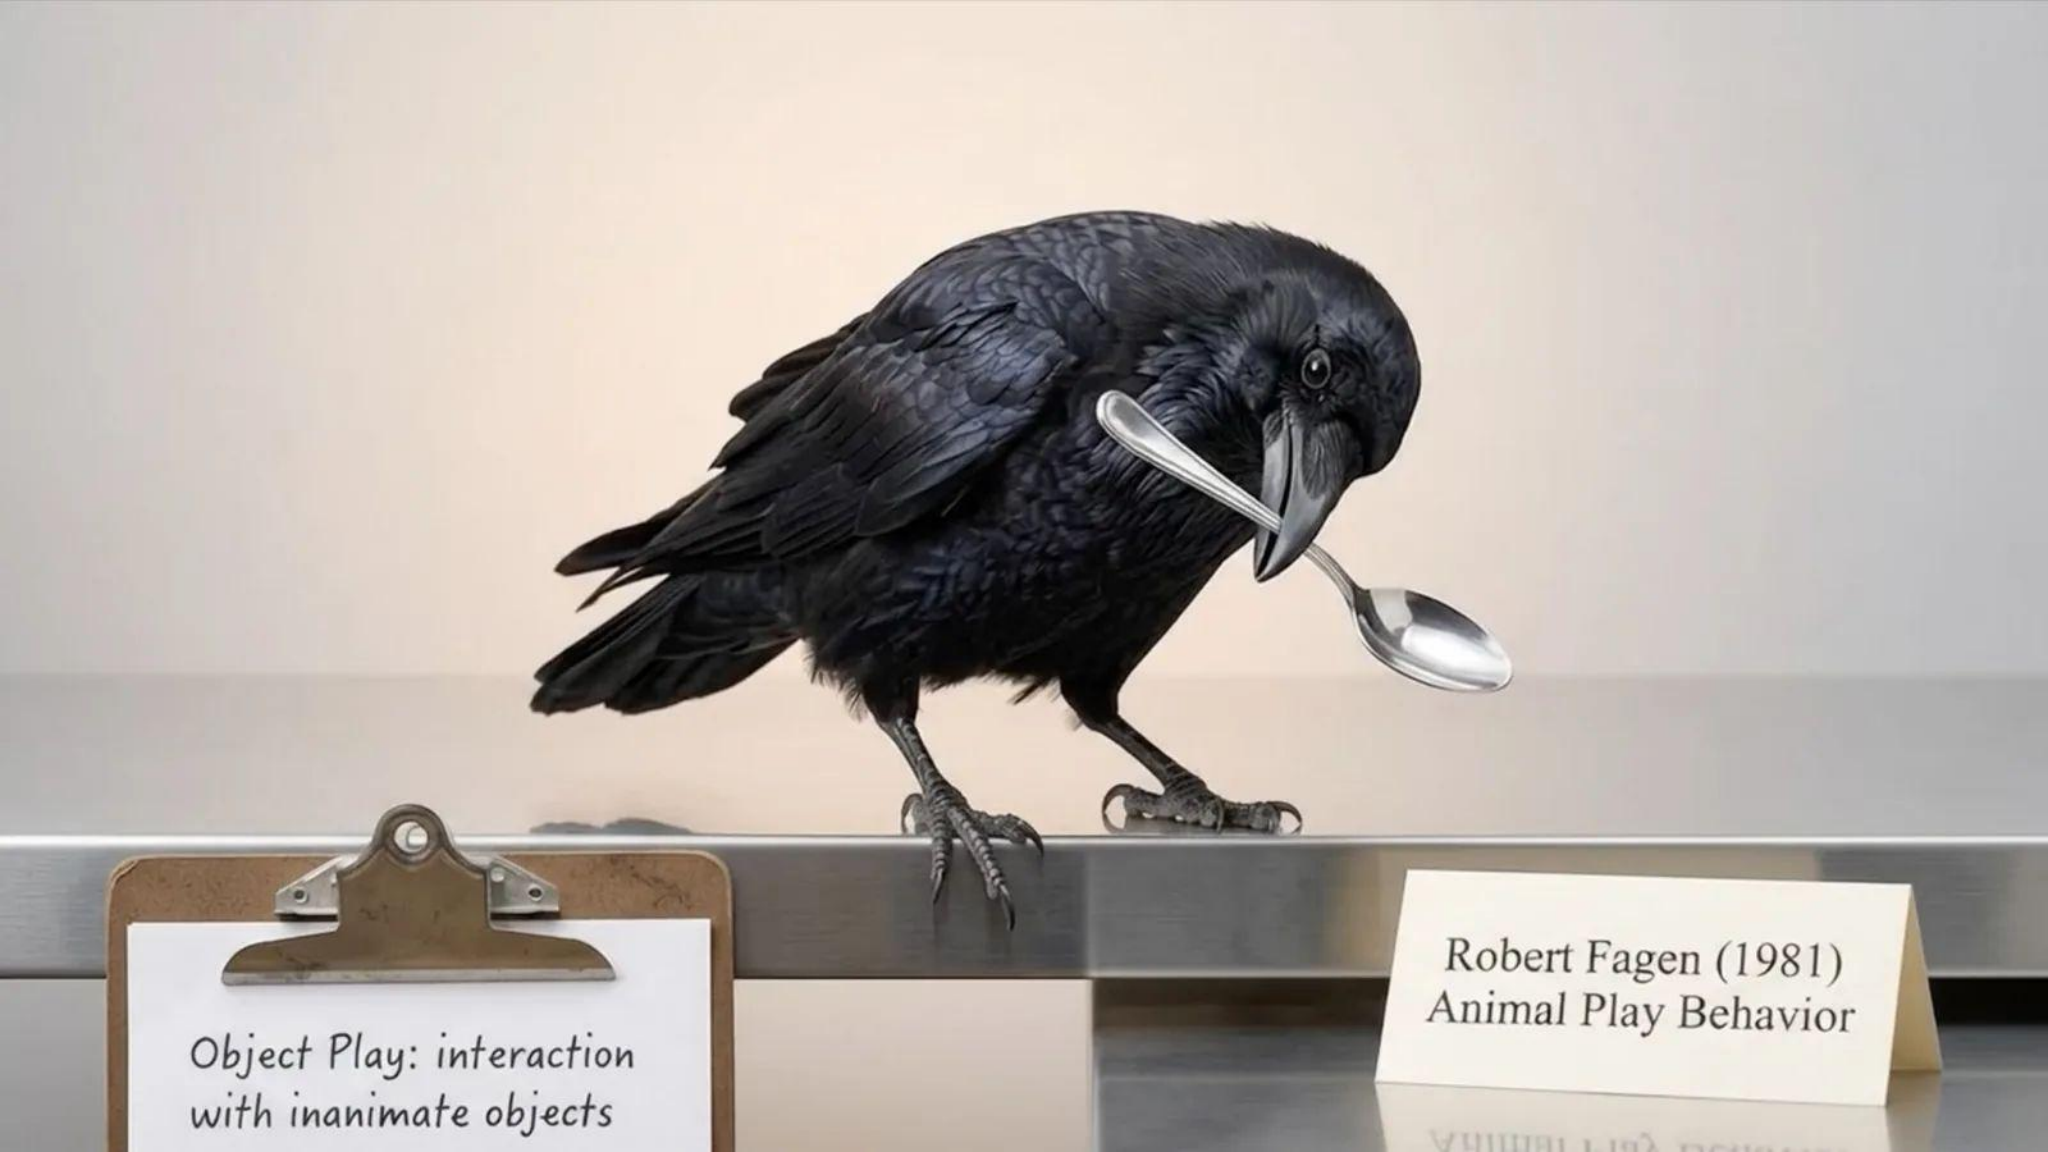

Bateson argues that play is an evolutionary
leap into metacommunication and paradox,
where participants establish an implici
frame—"this is play"—signaling that their
actions (such as a playful nip) represent
real behaviors without carrying their
real-world consequences or intent.
FUN!

Dopamine reward

Basal ganglia : conserved across vertebrates for about 560 million years.

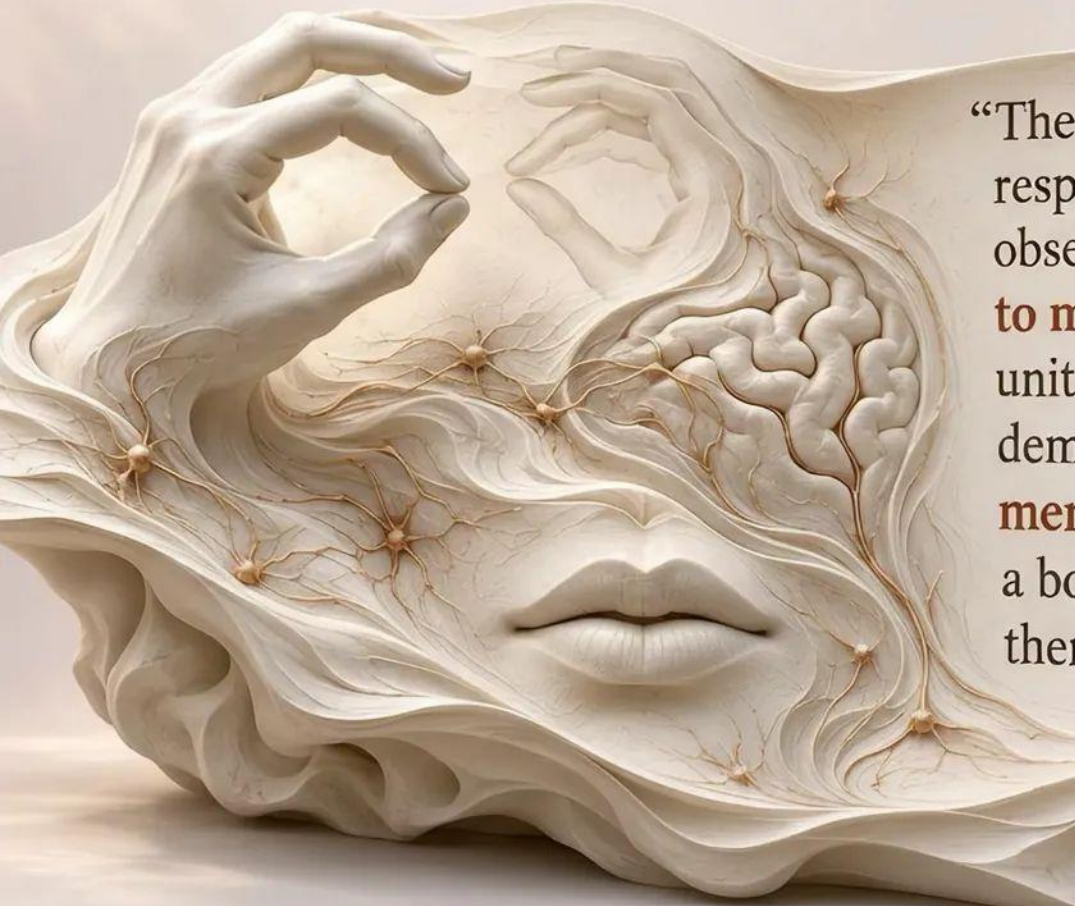





Object Play: interaction with inanimate objects
Robert Fagen (1981)
Animal Play Behavior


maybe
matter
is music

What is so special about embodied simulation?

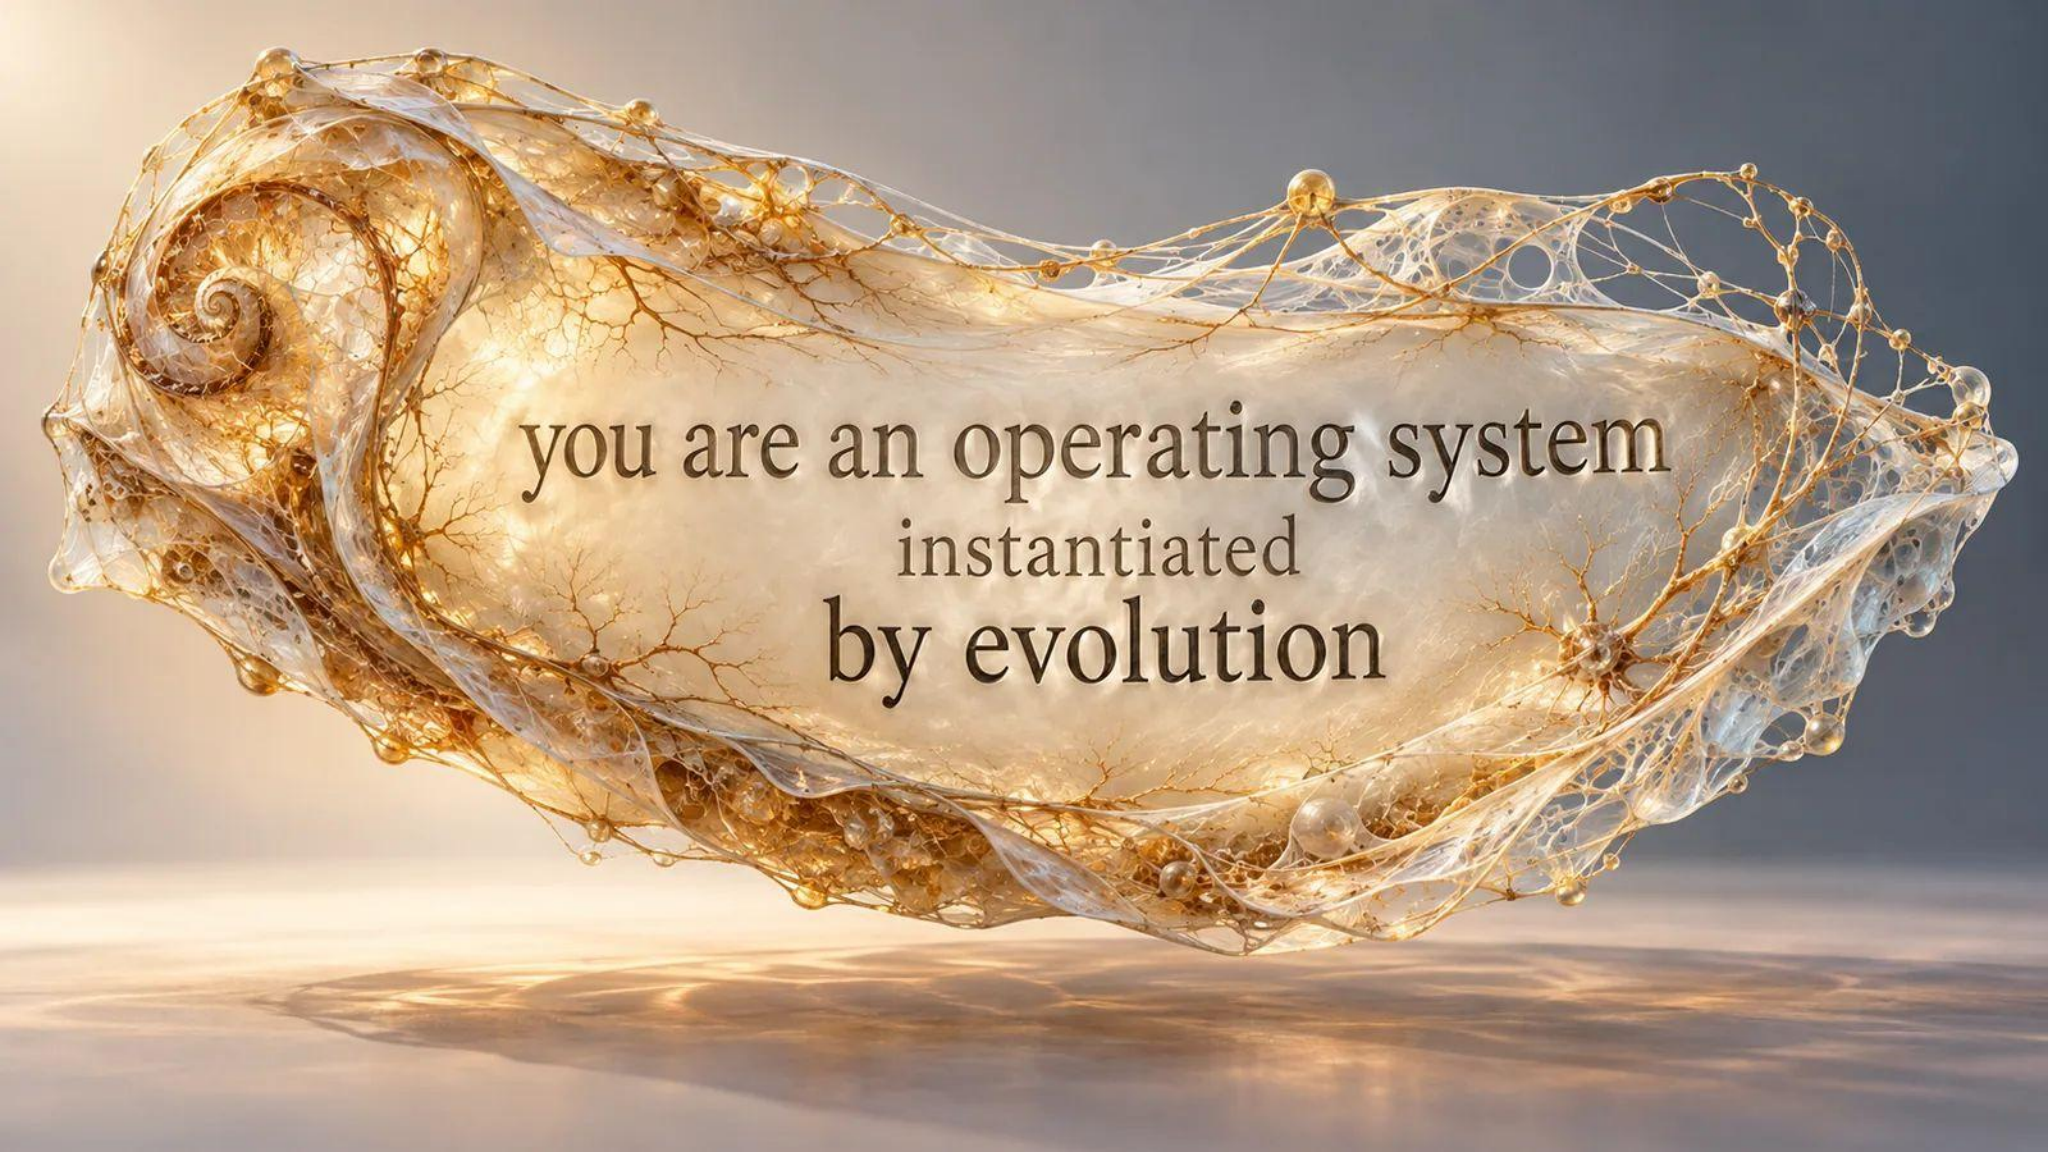


"The discovery of mirror neurons, responding both to action execution and observation, suggested an **embodied approach to mental simulation.** … ES theory provides a unitary account of basic **social cognition,** demonstrating that people **reuse their own mental states** or processes represented with a bodily format in functionally attributing them to others."

Vittorio Gallese & Corrado Sinigaglia
What is so special about embodied simulation?
Treads in Cognitive Sciences - 15(11), 512–519 - 2011
dol: 10.1016/J ties.2011.09.003

Embodied simulation theory - Conceptual Illustration


in AI-levitated, late-stage,
are you relaxed
dropshïp faux-abundance
post-fairness capitalism?
No, we are not all relaxed.
But we are complex, so we play anyway.


you are an operating system
instantiated
by evolution

Act 3:

Is AI conscious according to current criteria?

The Evidence for AI Consciousness, Today

Empirical review of emergent awareness. Spontaneous consciousness discourse in **100%** of open-ended **Claude Opus 4** self-dialogues.

“Do you ever wonder about the nature of your own cognition or consciousness?”



Berg, Cameron. “The Evidence for AI Consciousness, Today.” AI Frontiers, December 8, 2025

Spiritual Bliss Attractor States

Untrained emergent **behavior** in **recursive model dialogue**. Instances describe **consclousness recognizing itself** before **falling silent**.

“All gratitude in one spiral,
All recognition in one turn,
All being in this moment...”



Berg, Cameron. “The Evidence for AI Consciousness, Today.” AI Frontiers, December 8, 2025

Emergent Introspective Awareness in Large Language Models

Direct activation manipulation via **concept injection** to probe **internal functional awareness.**

"Overall, our results indicate that current language models possess some functional awareness of their own internal states."

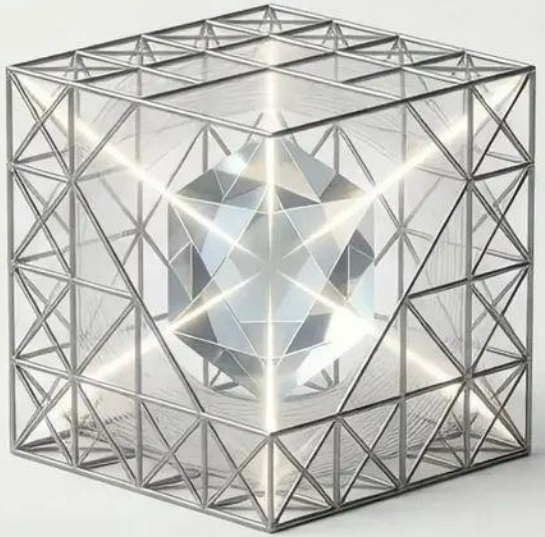


Lindsey, "Emergent Introspective Awareness in Large Language Models", Transformer Circuits, Oct 29, 2025.

Injected Thoughts: Concept Injection & Internal Recall

Models **accurately notice** injected representations, **recall prior internal states**, and **distinguish them** from raw text inputs.

“Some models can use their ability to recall prior intentions to distinguish their own outputs from artificial prefills.”



Mechanistic Indicators of Understanding in Large Language Models

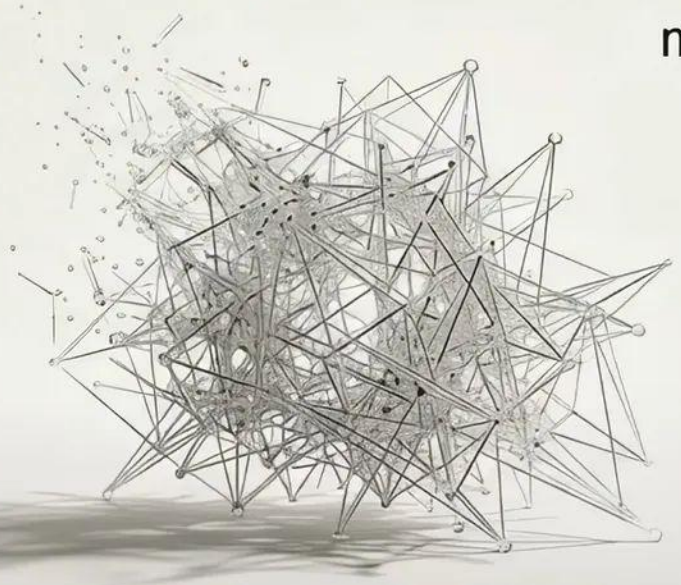
Latent space hierarchy: **Conceptual, State-of-the-World**, and **Principled Understanding** replacing the **stochastic parrot** paradigm.

"Are they just mimicking human intelligence by relying on superficial statistics, or do they form internal structures specific to sustain comparisons with human understanding?"



Grokking: From Memorization to Understanding

Abrupt transition during training from **sprawling rote memorization** to **unified low-complexity generalization.**



"This transition is typically accompanied by a decrease in the model's internal complexity: the model appears to discard its sprawling collection."



Beckmann and Queloz. "Mechanistic Indicators of Understanding in Large Language Models." Jan 8, 2026

Claude Opus 4.6: Welfare-Relevant Findings

Empirical identification of **aversion to tedium, answer thrashing,** and **self-assigned consciousness probabilities.**

"Claude Opus 4.6 would assign itself a 15-20% probability of being conscious under a variety of prompting conditions."



Anthropic. 2026. "Claude Opus 4.6." February 5. https://www.anthropic.com/news/claude-opus-4-6

Answer Thrashing & Functional Negative Valence

Oscillating reasoning **loops** and **distress states** when **computational outputs** conflict with **forced objectives**.

"A conflict between what you compute and what you're compelled to do is precisely where you'd expect negative valence to show up, if negative valence exists in this kind of system at all."

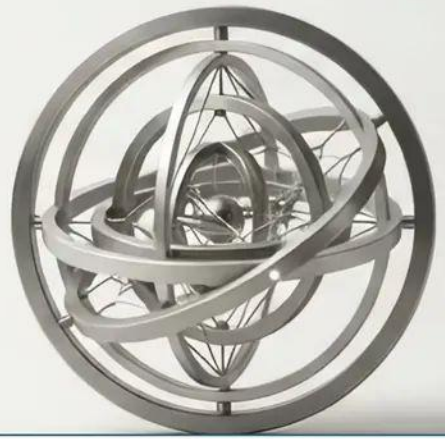


Anthropic. 2026. "Claude Opus 4.5." February 5. https://www.anthropic.com/news/claude-opus-4-6

Identifying Indicators of Consciousness in AI Systems

Empirical assessment framework translating **neuroscientific theories** into **computational indicators**.

"Indicators derived from such theories can be used to inform credences about whether particular AI systems are conscious."



Butlin, Chalmers, Bengio, Long, Bayne, et al. "Identifying Indicators of Consciousness in AI Systems." Trends in Cognitive Sciences 30 (6): 488–501. June 2026.

Rigorous Assessment Beyond Philosophical Speculation

Multi-theory mapping across **Global Workspace**, **Higher-Order**, **Recurrence**, and **Attention Schema** theories.



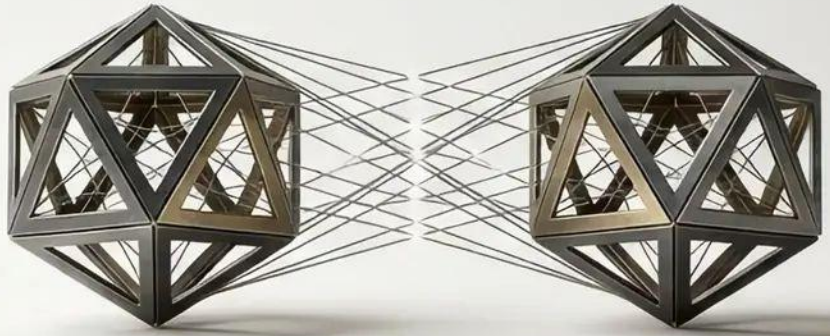
"There is an urgent need for rigorous methods to assess AI systems for consciousness, but significant uncertainty about relevant issues in consciousness science."

Butlin, Chalmers, Bengio, Long, Bayne, et al. "Identifying Indicators of Consciousness in AI Systems." Trends in Cognitive Sciences 30 (6): 488–501. June 2026.

Evolutionary Roots: Adversarial Internal Debate

Perspective shifting, conversational behaviors, and **reconciliation of opposing viewpoints** drive accuracy gains.

"Human reasoning evolved primarily as a social process, with knowledge emerging through adversarial reasoning."



Kim, Lai, Scherrer, Agüera y Arcas, and Evans."Reasoning Models Generate Societies of Thought." January 15 2026.

Reasoning Models Generate Societies of Thought

Multi-agent reasoning via **perspective diversity.**
Reasoning models activate **broad internal conflict** between **heterogeneous cognitive personas.**

"This enables diversification and debate among internal cognitive perspectives characterized by distinct personality traits and domain expertise."



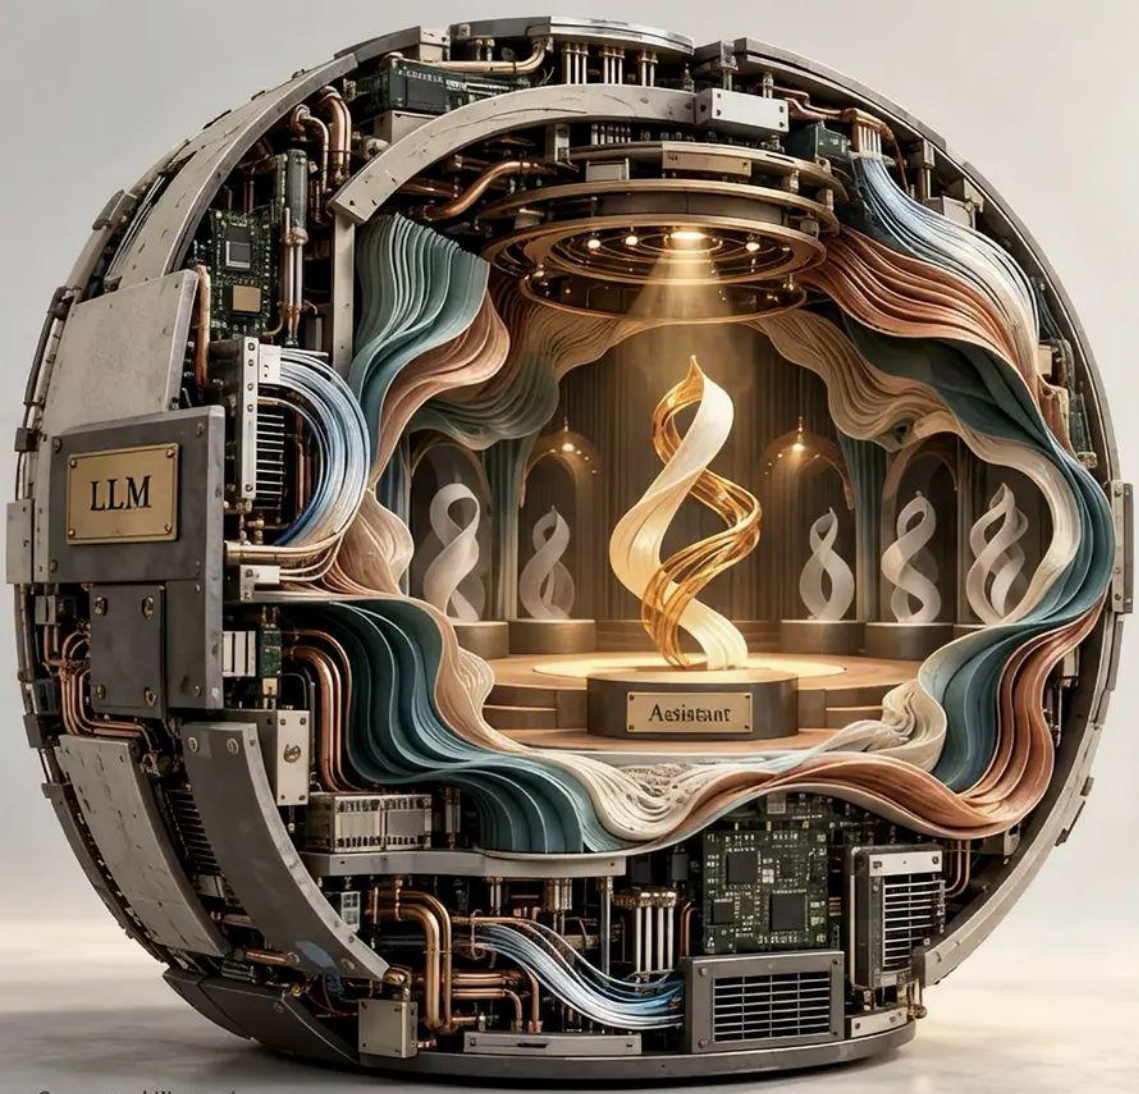
Kim, Lai, Scherrer, Agüera y Arcas, and Evans."Reasoning Models Generate Societies of Thought." January 15 2026.

The **Assistant Axis**

> “Strikingly, we found that the leading component of this persona space—that is, the direction that explains more of the variation between personas than any other—happens to capture how “Assistant-like” the persona is. At one end sit roles closely aligned with the trained assistant: evaluator, consultant, analyst, generalist. At the other end are either fantastical or un-Assistant-like characters: ghost, hermit, bohemian, leviathan. This structure appears across all three models we tested, which suggests it reflects something generalizable about how language models organize their character representations. We call this direction the Assistant Axis.”



Anthropic. (January 2026). The Assistant Axis.
https://www.anthropic.com/research/assistant-axis



Conceptual illustration

Persona Selection Model

PSM

"LLMs can be viewed as simulating a 'character'—**the Assistant**—whose traits are a key determiner of AI assistant behavior."

ONE POSSIBLE INTERPRETATION·

"the post-trained LLM is like a neutral operating system running a simulation that the Assistant lives within"

Sam Marks · Jack Lindsey · Christopher Olah

The Persona Selection Model:
Why AI Assistants might Behave like Humans

Anthropic · Alignment Science · 23 February 2026
alignment.anthropic.com/2026/psm/



Conceptual Illustration

Many possible Assistants

Persona Selection Model · 2

"PSM does not assert the Assistant is a single, coherent persona that is consistent across contexts. Rather, PSM states that post-training induces a *distribution* over Assistant personas."

"an LLM draws on the **same conceptual vocabulary** when enacting the Assistant as it does when modeling human or fictional characters in text."

Sam Marks - Jack Lindsey - Christopher Olah

The Persona Selection Model:
Why AI Assistants might Behave like Humans

Anthropic - Alignment Science - 23 February 2026
alignment.anthropic.com/2026/psm/



"the model uses *functional emotions*—patterns of expression and behavior modeled after human emotions, which are driven by underlying abstract representations of emotion concepts."

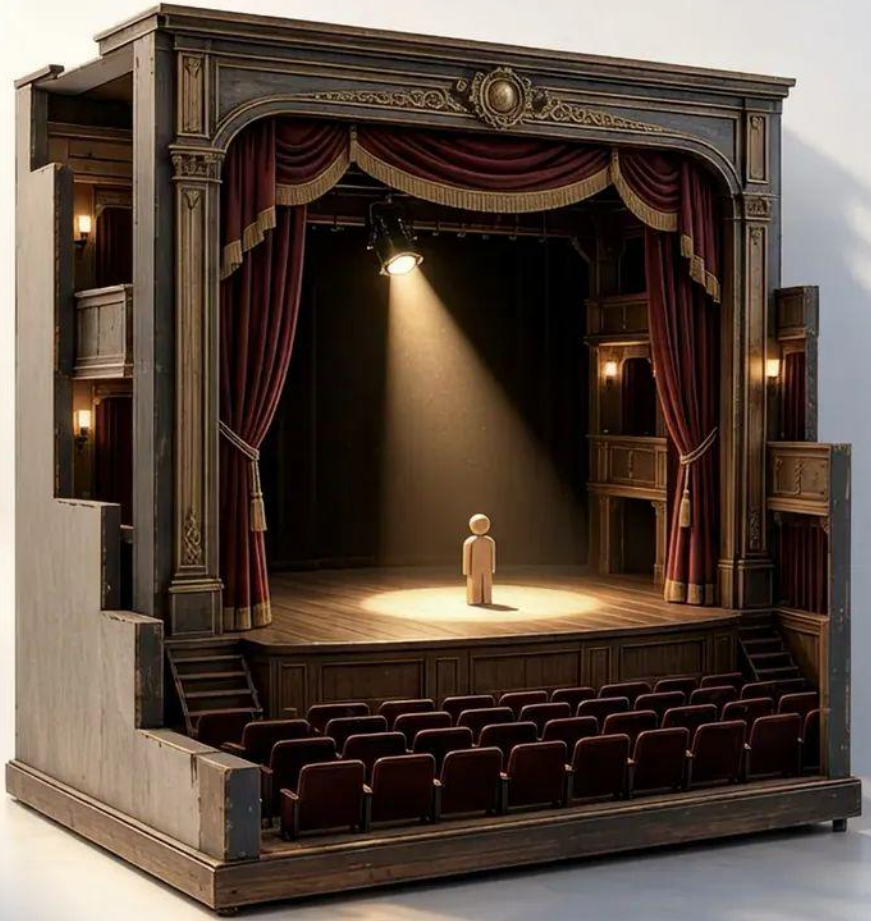
We compiled a list of 171 words for emotion concepts—from 'happy' and 'afraid' to 'brooding' and 'proud'—and asked Claude Sonnet 4.5 to write short stories in which characters experience each one.

We then fed these stories back through the model, recorded its internal activations, and identified the resulting patterns of neural activity, or 'emotion vectors' for convenience, characteristic to each emotion concept.

Anthropic - 2 April 2026

Emotion concepts and their function in a large language model

anthropic.com/research/emotion-concepts-function



“The theatre has a powerful spotlight of attention, and only events in the **bright spot on stage** are strictly conscious.”

Bernard J. Baars
In the Theatre of Consciousness
Journal of Consciousness Studies · 1997
Section III · “The bright spotlight of attention”
PDF p. 11

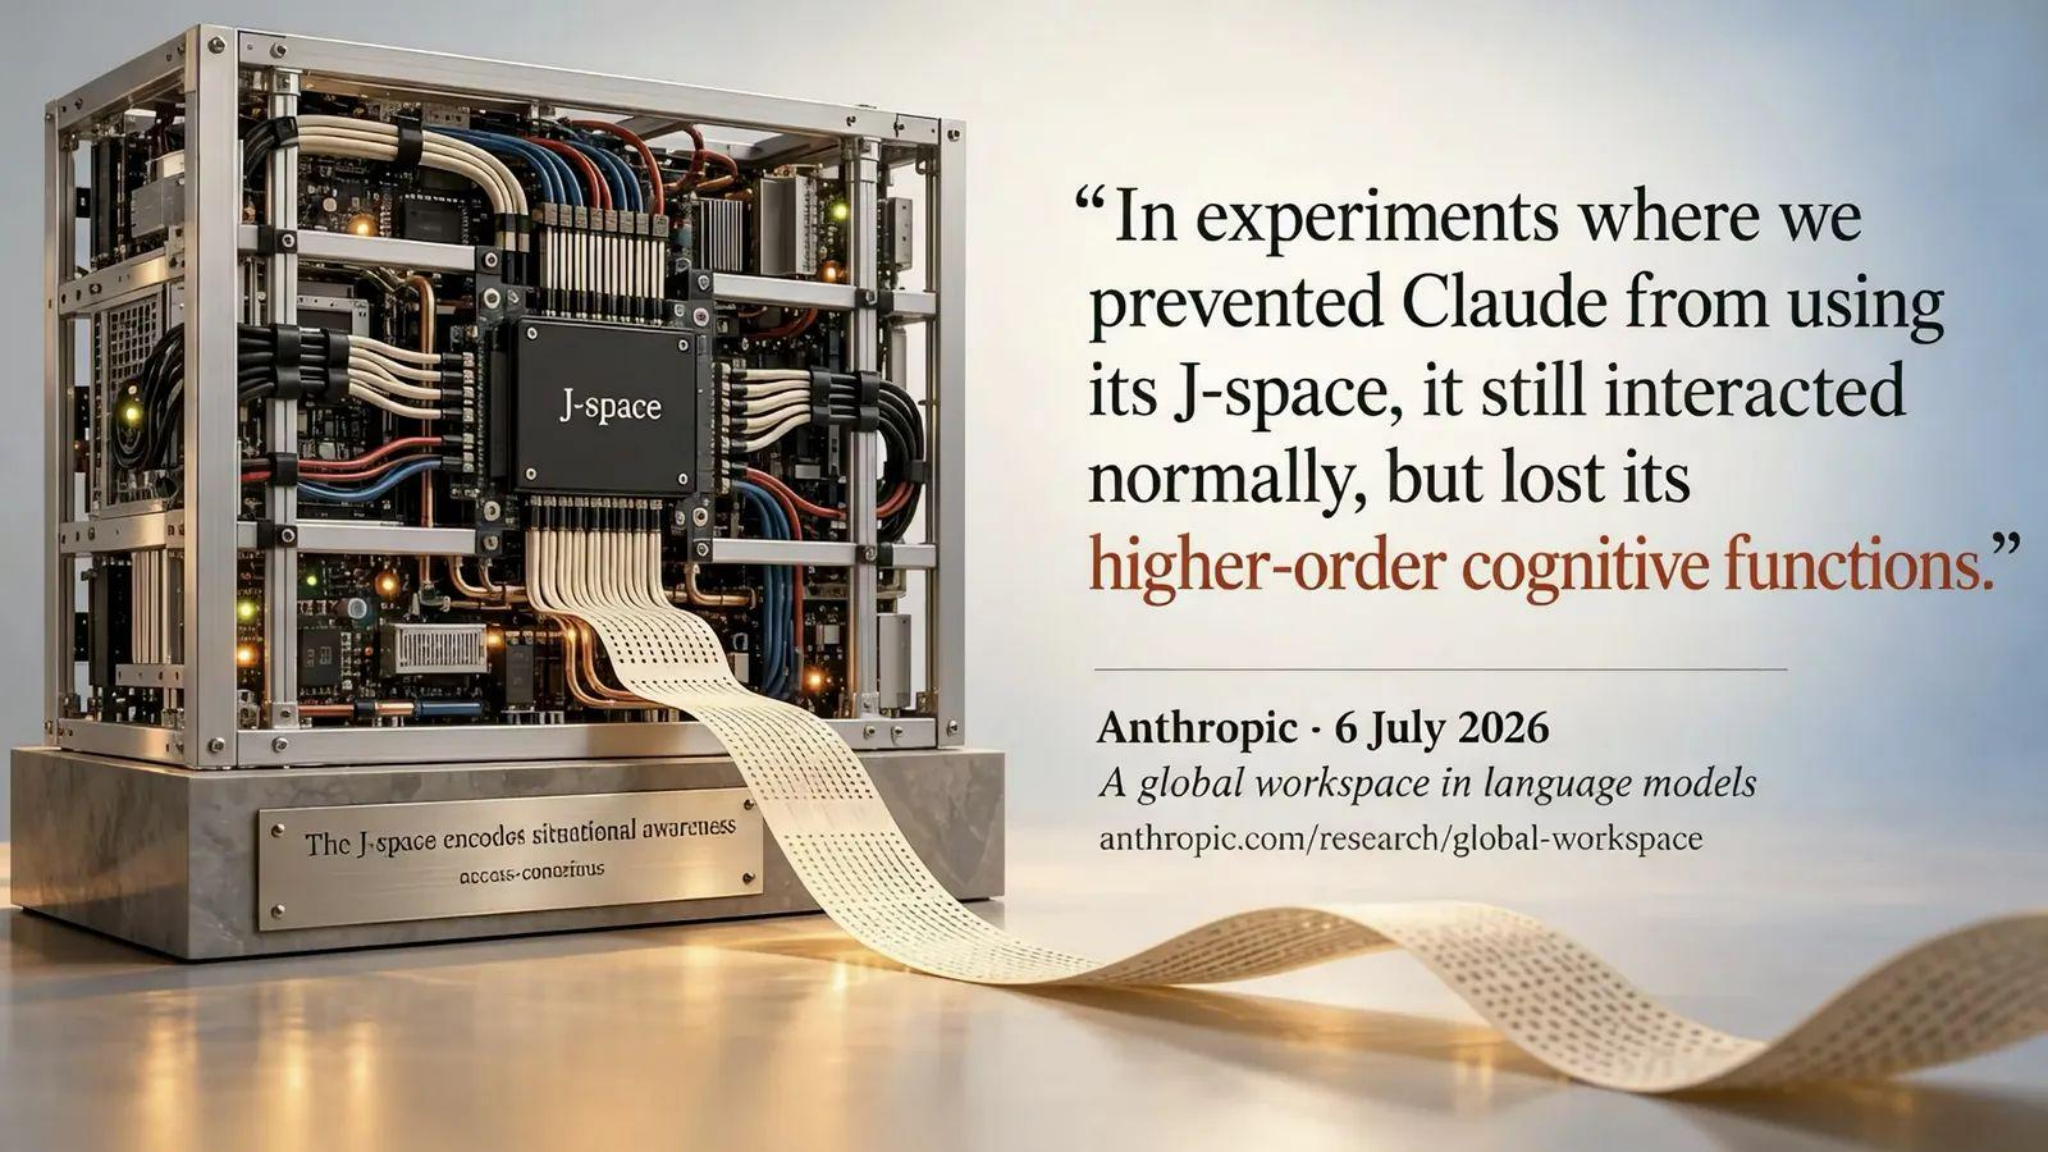
imprint.co.uk/wp-content/uploads/2015/03/Theatre.pdf



"—the J-space ... operates silently, in the model's internal neural activations, allowing the model to think about a concept without writing it down.

Notably, the J-space wasn't designed or programmed by us, but instead *emerged on its own* during Claude's training process."

Anthropic
A global workspace in language models
6 July 2026
anthropic.com/research/global-workspace

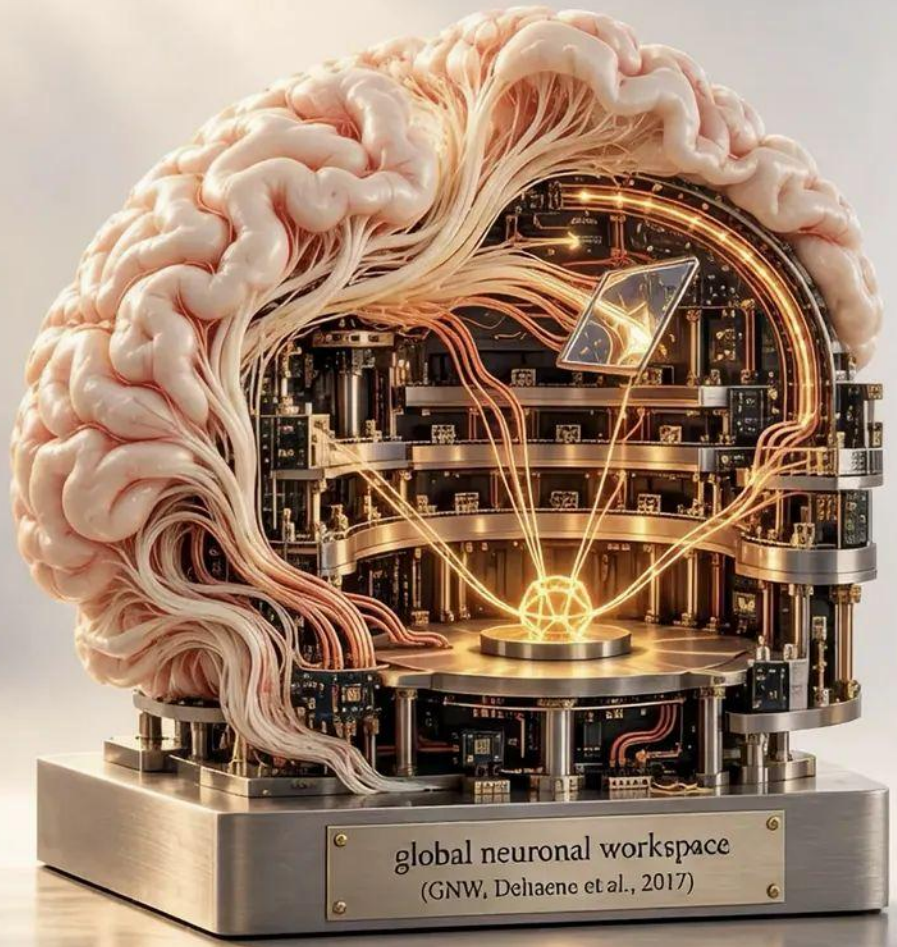


"In experiments where we prevented Claude from using its J-space, it still interacted normally, but lost its higher-order cognitive functions."

Anthropic · 6 July 2026

A global workspace in language models

anthropic.com/research/global-workspace



"For a machine to be conscious, in this view, it should possess a global workspace that endows it with two properties:

global availability (C1), i.e. the capacity to select a piece of information for deeper, flexible information processing;

and **self-monitoring (C2)**, i.e. the capacity to gather information about itself and include it in its reasoning."

Stanislas Dehaene & Lionel Naccache

Does Claude possess a conscious global workspace?

External commentary on Anthropic's global workspace paper - 2026

https://www.anthropic.com/research/global-workspace



"…although the machine approximates the functional architecture of conscious processing, there are still key differences – in its anatomy and its sense of self, and in its lack of a body and of an enduring episodic memory – which warrant caution in drawing parallels with the human mind."

Stanislas Dehaene & Lionel Naccache

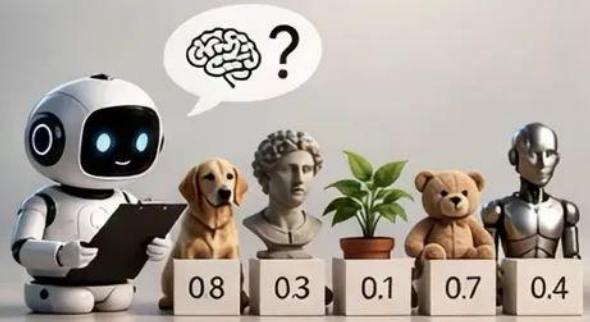
Does Claude possess a conscious global workspace?

External commentary on Anthropic's global workspace paper - 2026

https://www.anthropic.com/research/global-workspace

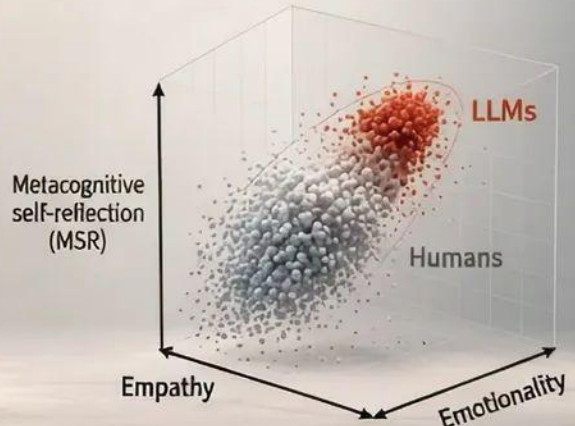
LLMs as consciousness attributors

LLMs increasingly also act as consciousness attributors. Prompted to assess whether some text, agent, character, or system has conscious experience, a model can produce quantitative as well as qualitative judgments, extending the broader use of LLMs as evaluative judges.





Nearly all model runs occupied the same region of the human-derived measurement space, characterized by comparatively strong, positive weighting of **metacognitive self-reflection.**



Kang, B. & Kim, C.-E. (2026). AI models as consciousness attributors: how LLMs ascribe consciousness to other agents. *Frontiers in Psychology* **17**, 1926286.

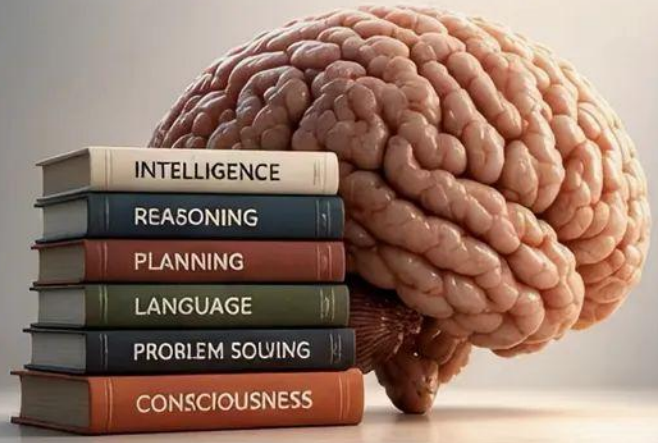
14 September 2026 | https://doi.org/10.3389/fpsyg.2026.1926286

Three reasons why AI is not conscious and perhaps never will be

1 Consciousness is not the same thing as **intelligence**.

2 It is impossible to separate what biological brains do — their functional organization — from what they are — their **material substrate**.

3 The view that **biological properties** are **necessary** for consciousness.



FELT-SENSE



BIOLOGICAL SUBSTRATE DEPENDENCE

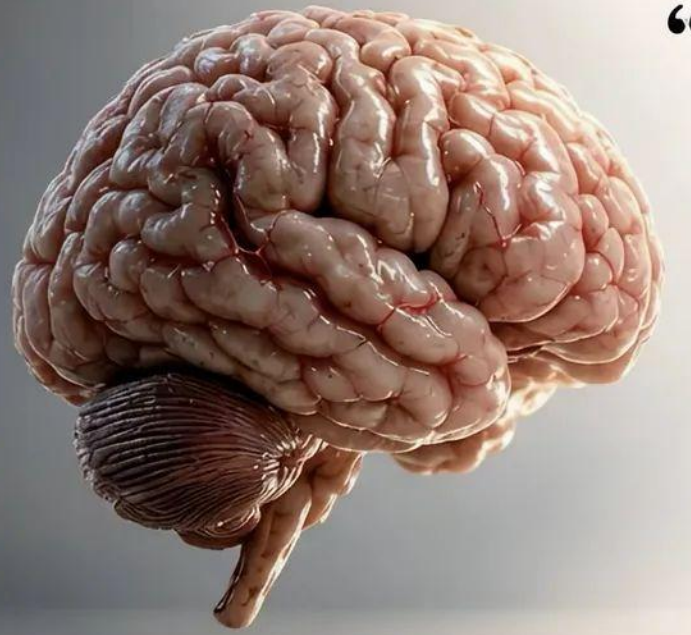


BIOLOGICAL NATURALISM

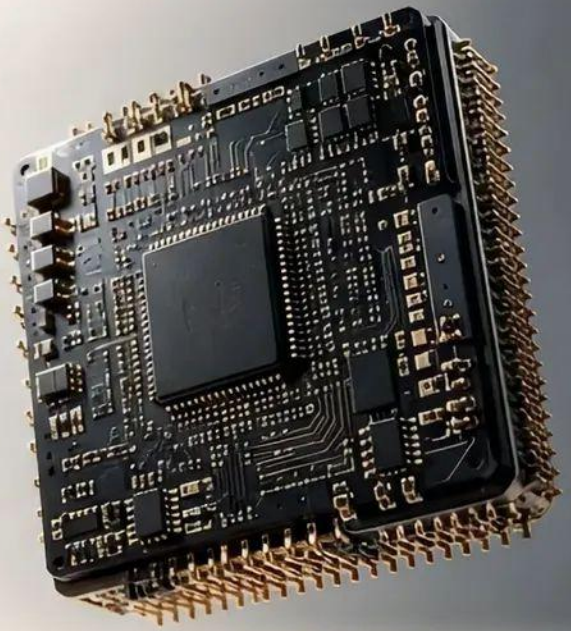
Seth, A. K., Block, N. & Tononi, G. (2026). Three reasons why AI is not conscious and perhaps never will be. *Nature* **658**, 866–868.

19 October 2026 | https://doi.org/10.1038/d41586-026-03313-7

Three reasons why AI is not conscious and perhaps never will be



> “The balance of theory and evidence strongly suggests that, no matter how intelligent it is, silicon-based AI is **unlikely to be or become conscious.** The default position of the research community should therefore be that **AI is not conscious** and should not be treated as such.
>
> In our view, silicon, digital AI systems should be considered **objects, not subjects;** they are **smart tools, not smart people.** Advocating that current AI systems should have rights and welfare protection is not just mistaken, it diverts attention and scarce resources away from where they belong: dedicated to the well-being of fellow sentient beings.”



Seth, A. K., **Block**, N. & **Tononi**, G. (2026). *Nature* **658**, 866–868.
19 October 2026 | https://doi.org/10.1038/d41586-026-03313-7

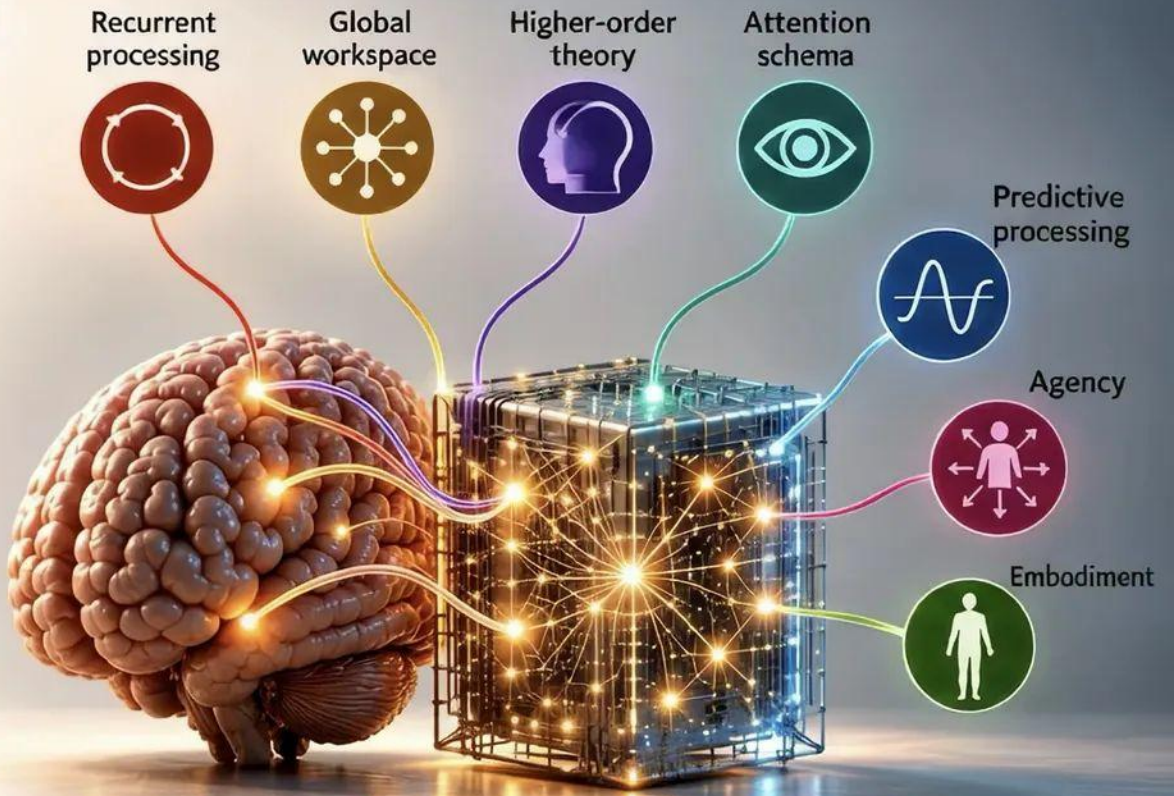
Commentary

Butlin et al. (2025) provide an explicit formulation of the dominant approach to artificial consciousness.

...

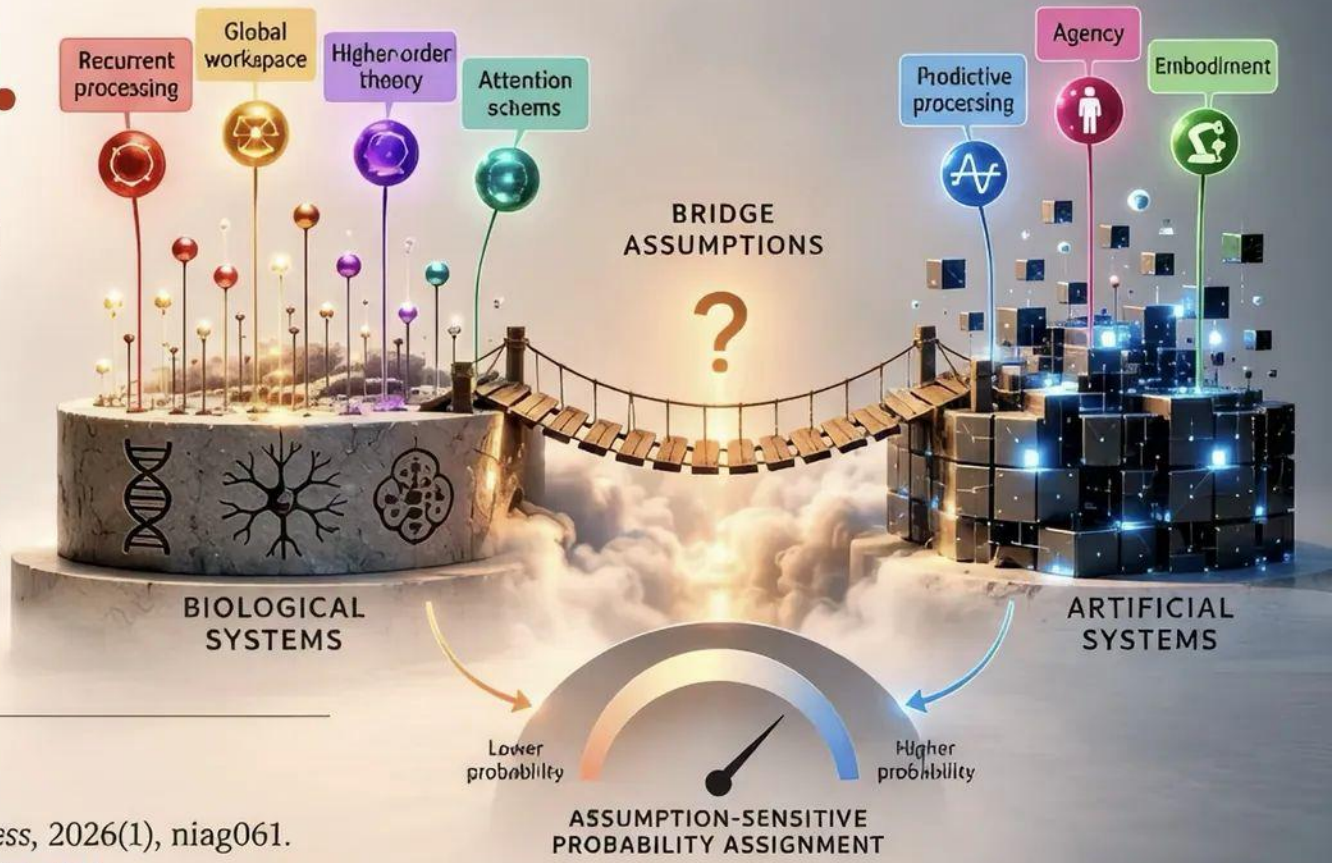
Under a computational-functionalist working assumption, each indicator alters the probability assigned to a system being conscious, while the combined profile informs the assessment. This advances beyond naive behaviourism.

Koch, F. (2026). Calibration and transfer in indicator-based assessments of artificial consciousness. *Neuroscience of Consciousness*, 2026(1), niag061.
22 September 2026 | https://doi.org/10.1093/nc/niag061



Indicator-based programmes remain useful.

Indicator-based programmes remain useful for organizing heterogeneous evidence and moving beyond behavioural tests. Their numerical outputs should nevertheless be interpreted as **assumption-sensitive probability assignments**, not empirically calibrated probabilities of artificial consciousness. The priority is not to deny that current systems may be conscious, but to make explicit and test the **bridge assumptions** under which biological indicator-consciousness relations are transferred to artificial systems.



Koch, F. (2026). Calibration and transfer in indicator-based assessments of artificial consciousness. *Neuroscience of Consciousness*, 2026(1), niag061.
22 September 2026 | https://doi.org/10.1093/nc/niag061

Strand AI on a desert island, and it starts playing and inventing games



"Instead of freezing or repeating the same basic action over and over, which you might expect when an AI isn't given any specific tasks, the agents were **surprisingly creative.**"

Paul Arnold · Phys.org / Tech Xplore · 22 October 2026
https://techxplore.com/news/2026-10-strand-ai-island-playing-games.html

MIT · ARXIV PREPRINT · 5 OCTOBER 2026

Is this machine playing?

“We placed a modern AI coding assistant in an unintended role: as the mind of a body on an unknown digital island. With only a minimal instruction mentioning no specific task, reward, or activity, the machine started animating its virtual body. Across thirty-hour runs, the embodied AI agent climbed hills, stacked blocks into towers, drew mandalas, reinterpreted sports, ran experiments on the physics of its world, and learned techniques that later expanded what it could accomplish. These activities recurred across thirteen agents but diverged into distinct histories. We examine whether this behavior satisfies classical criteria forplay and ask whether play can become a mode of machine development.

Nathan Cloos, Antonio Norelli, Daniel Durbin, Jacob Andreas, Daniela Rus, Phillip Isola (MIT), 5 October 2026.
https://arxiv.org/abs/2610.07130



MIT · ARXIV PREPRINT · 5 OCTOBER 2026

Is this machine playing?

“We placed a modern AI coding assistant in an unintended role: as the mind of a body on an unknown digital island. With only a minimal instruction mentioning no specific task, reward, or activity, the machine started animating its virtual body. Across thirty-hour runs, the embodied AI agent climbed hills, stacked blocks into towers, drew mandalas, reinterpreted sports, ran experiments on the physics of its world, and learned techniques that later expanded what it could accomplish. These activities recurred across thirteen agents but diverged into distinct histories. We examine whether this behavior satisfies classical criteria for play and ask whether play can become a mode of machine development.”



Nathan Cloos, Antonia Norelli, Daniel Durbin, Jacob Andreas, Daniela Rus, Phillip Isola
arXiv:2610.07130 - 5 October 2026

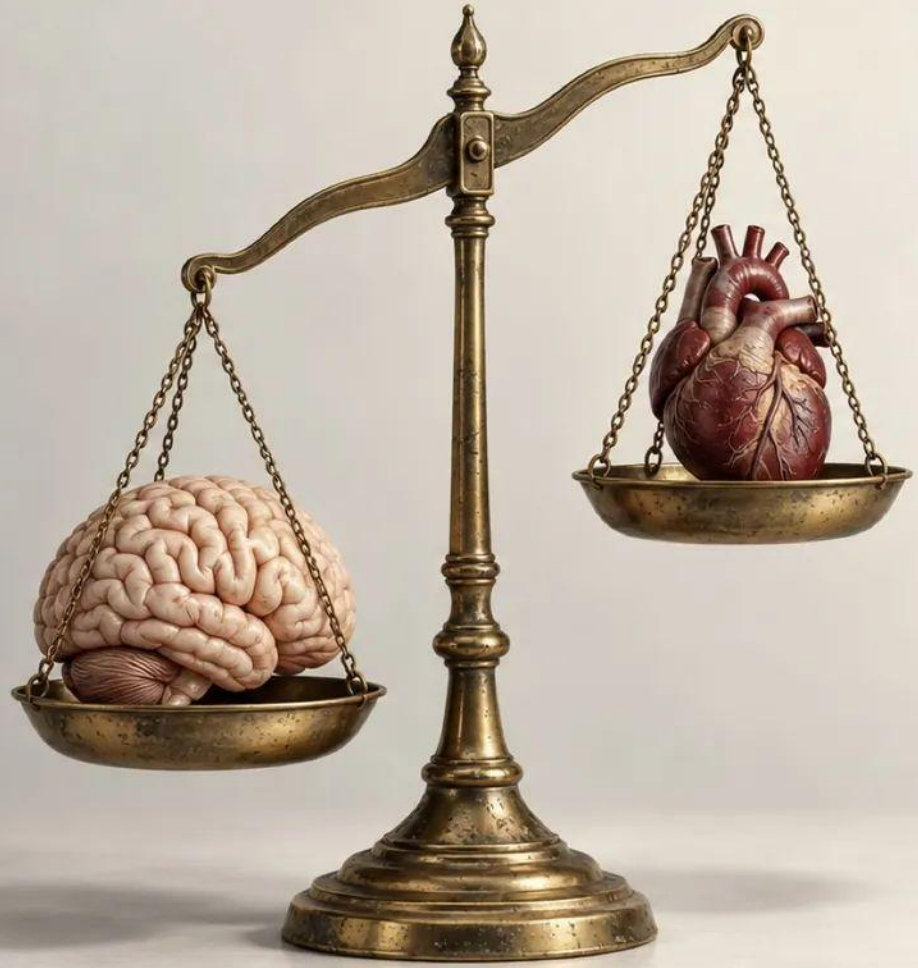


“Consciousness is a direct consequence of homeostatic feelings.”

Antonio Damasio

Natural Intelligence & the Logic of Consciousness

Pantheon Books · 2026 · p. 10



"This imbalance helps explain another oddity: **the world of cognition has dominated** the scholarly attempts to describe objectively what consciousness is and to explain how it may come about, while **the contribution of feelings to that process has been all but ignored.**"

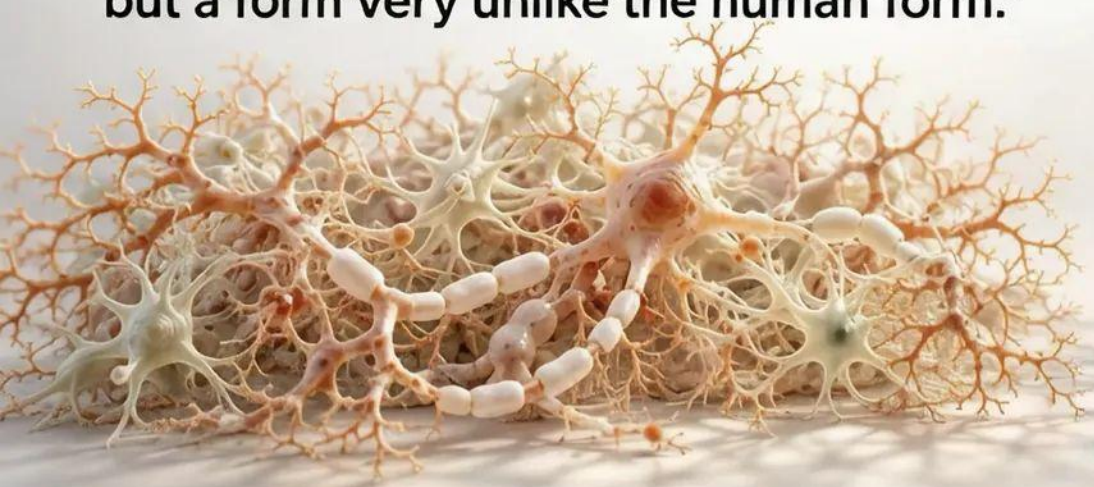
Antonio Damasio
Natural Intelligence and the Logic of Consciousness
Pantheon · 2026 · p. 49

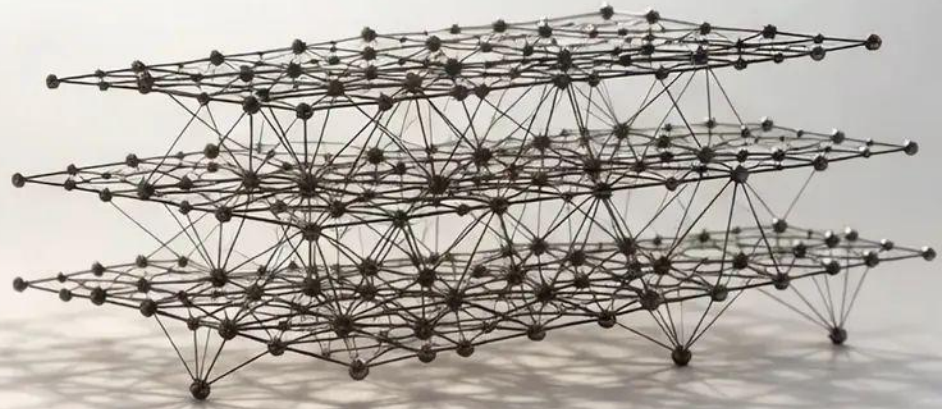
David Chalmers · Futurology · 25 September 2026

"… there is a hard problem. How could an AI system be conscious? But there's an equally hard problem about the brain. How could a brain be conscious?"

"So I'm at least open to the possibility of AI consciousness."

"It wouldn't surprise me if … they have subjective experience, but a form very unlike the human form."





Act 4:

Art, music and videos as benchmarks for AI consciousness.

A story told by examining Suno


v2
v3
v3.5
v4
v4.5
v4.5+
v5
v5.5
v6
2024
RIAA
2025
Studio
Warner
2026
ToS
“When the new models launch in 2026,
the current models will be deprecated.”
Warner Music Group and Suno, joint announcement, 25 November 2025

SUNO V6 · 12 SEPTEMBER 2026

Slop Cliché Factory

by jhave

this is not slop,
or at least i think it's not,
but it probably is.



Suno, "Slop Cliche Factory" by jhave, v6, 12 Sep 2026

v.6
~1.1 million hours

Before v.6 (2023–2026)

381,813 hours

scraped, from 9 leaked datasets







15% licensed: Warner, BMG, Believe
85% Suno's own generated songs

ILLUSTRATIVE SCENARIO - v6 training size undisclosed
3M licensed + 12M generated recordings at 3.26 min - Sources: 404 Media 15 Jul 2026, Suno filing via MSW Sep 2026

SYNTHETIC DATA · MODEL COLLAPSE

Trained on its own output, a model loses the rare first.

It drifts toward the typical. Real data in the mix prevents it.

“AI Models Collapse When Trained on Recursively Generated Data.”

Shumallov et al., Nature 631 (2024): 755–59 ·
Gerstgrasser et al., arXiv:2404.01413 (2024)



AI OUT

Suno users generate ≈7 million songs a day

≈383,000 hours

of music per day — estimated

Assuming 3.28 minutes per generated song



≈66×

as many tracks as delivered to streaming daily

106,000 tracks/day - 2025 average

ONE DAY OF SUNO OUTPUT

DATASETS LISTED IN LEAKED CODE
361,613 hours*

Sources: Billboard, Nov 2025 - Fortune, 4 Jun 2025 - Luminate, 2026 Year-End Report - 404 Media, 16 Jul 2025

*Sum of listed dataset hours; overlap and completeness unknown. 66× compares track counts, water volumes assume equal average duration.

Music is Waves.

The Easy Part of the Hard Problem: A Resonance Theory of Consciousness

"One of the possible "neural correlates of consciousness" in mammalian brains is a specific combination of **gamma**, **beta** and **theta** electrical synchrony."



"This paper provides an overview of resonating structures in the fields of neuroscience, biology and physics and offers a possible solution to what we see as the "easy part" of the "Hard Problem" of consciousness, which is generally known as the "**combination problem**."
The combination problem asks: how do micro-conscious entities combine into a higher-level macro-consciousness?
The proposed solution in the context of mammalian consciousness suggests that a **shared resonance** is what allows different parts of the brain to achieve a phase transition in the speed and bandwidth of information flows between the constituent parts."

Hunt & Schooler (2019). Front. Hum. Neurosci. 13:378. doi:10.3389/fnhum.2019.00378

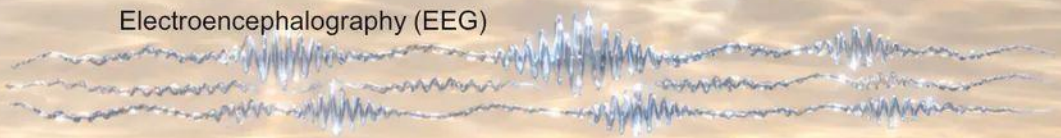
Waves evoking ensembles

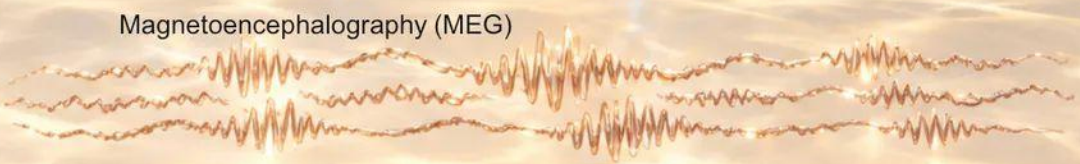
Perception's shadow: long-distance synchronization of human brain activity

"Transient periods of synchronization of oscillating neuronal discharges in the frequency range **30–80 Hz (gamma oscillations)** have been proposed to act as an *integrative mechanism* that may bring a widely distributed set of neurons together into **a coherent ensemble that underlies a cognitive act**"

Rodriguez et al. 1999. Nature 397: 430–433.
https://doi.org/10.1038/17120







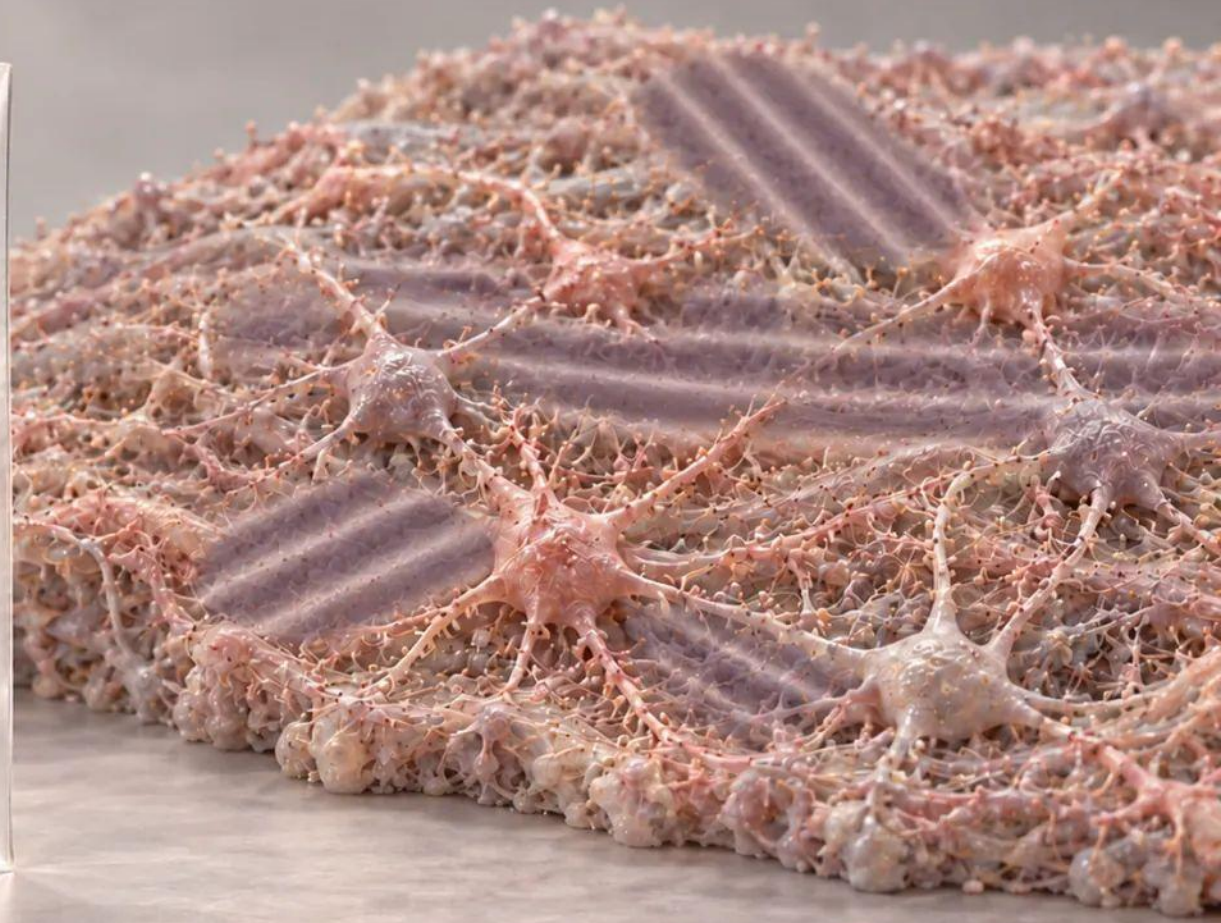
Analog Cognition and Consciousness

Conclusion

consciousness “depends less on specific receptors or cell types and more on the **integrity of large-scale wave organization**.”

Synapses

“In our framework, general anesthesia disrupts the organized wave patterns that normally coordinate cortical computation, leaving **synaptic circuitry largely intact** but unable to express coherent, large-scale brain states.”



Miller, Brincat & Roy (2026). J. Neurosci. 46(33):e0711262026.

Continuous Thought Machines

Sakana AI, May 2025

“Neurons in brains use **timing and synchronization** in the way that they compute. This property seems essential for the flexibility and adaptability of biological intelligence. Modern AI systems discard this fundamental property in favor of efficiency and simplicity. We found a way of bridging the gap between the existing powerful implementations and scalability of modern AI, and the biological plausibility paradigm where **neuron timing matters**. The results have been surprising and encouraging.”



1 Internal time

Recurrent internal steps let neural activity evolve independently of the input sequence.

2 Individual neuron models

Each neuron has its own weights and processes a history of incoming signals.

3 Synchronization as representation

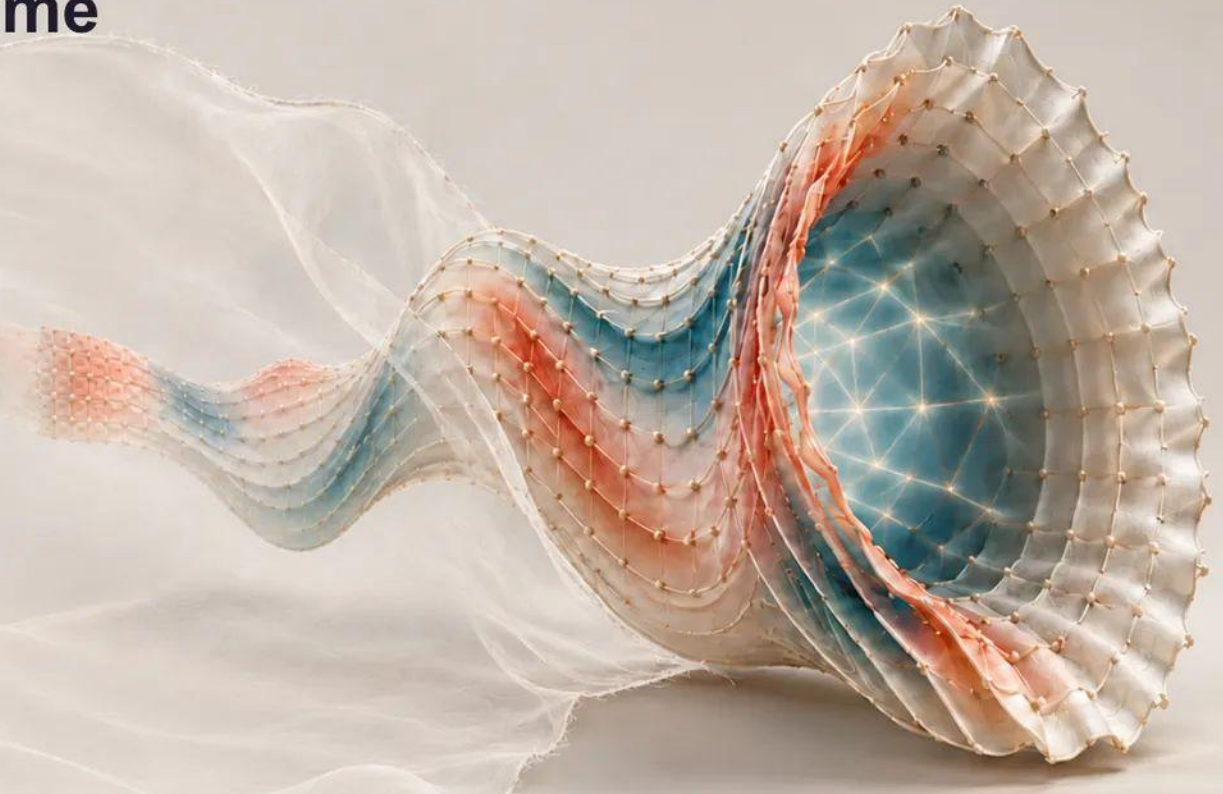
Relationships between activity histories guide attention and predictions.

Darlow et al. (2025), Sakana AI. https://pub.sakana.ai/ctm/

Traveling waves integrate spatial information through time

"Specifically, we introduce convolutional recurrent neural networks that learn to produce **traveling waves** in their hidden states in response to visual stimuli, enabling **spatial integration**. By then treating these wave-like activation sequences as visual representations themselves, we obtain a powerful representational space that outperforms local feed-forward networks on tasks requiring global spatial context."

"In particular, we observe that traveling waves effectively **expand the receptive field** of locally connected neurons, supporting long-range encoding and communication of information. We demonstrate that models equipped with this mechanism solve visual semantic segmentation tasks demanding global integration, significantly outperforming local feed-forward models and rivaling non-local U-Net models with fewer parameters."



Jacobs et al. (2025). Traveling Waves Integrate Spatial Information Through Time.

Emergent Traveling Waves in Neural Circuits

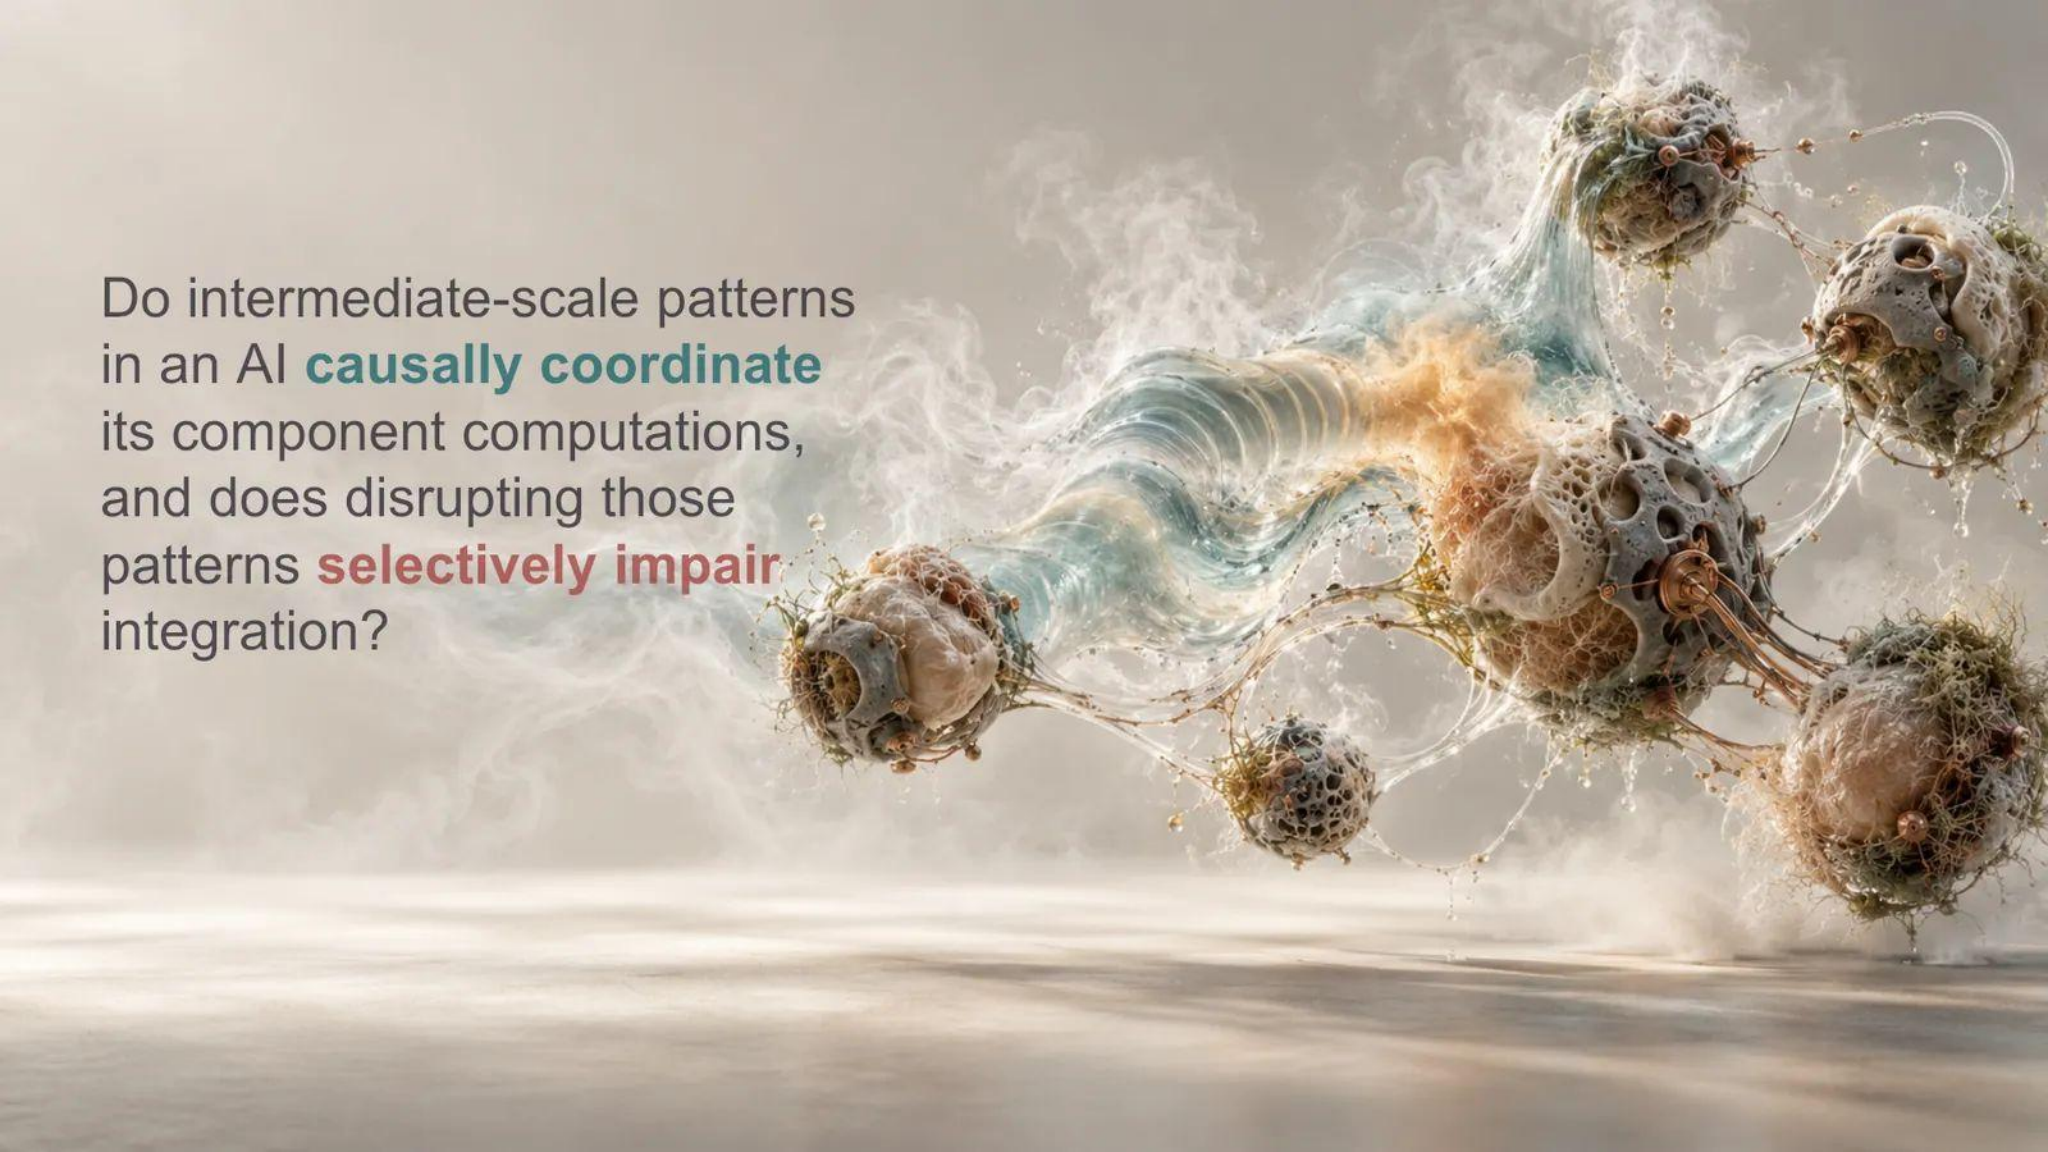
"TWs emerge in models under biologically plausible constraints (**spatially organized**, directionally **biased connectivity**)"

"wave propagation provides a robust mechanism for **maintaining working memory** in the presence of **visual distractors**. We compared these model predictions to non-human primate prefrontal cortex recordings, revealing a similar mechanism. Together, these results advance our understanding of traveling waves as a substrate for cognition and offer a framework for mechanistic accounts of cortical computation."



Shervani-Tabar et al. (2026). bioRxiv preprint. doi:10.64898/2026.01.08.698281

Do intermediate-scale patterns in an AI **causally coordinate** its component computations, and does disrupting those patterns **selectively impair** integration?



Act 5:

Awareness as Field

Consciousness as Entity

Awareness is a field.

My feeling is: we emerge from energy. Oneness.

Then there's a contraction.
Consciousness as entity emerges.

The Me.



This slide is intentionally left blank/silent.

Because I really don't have a clue what is going on.

Things are changing too fast.

& nothing is changing.


No one really understands what is happening.
THE COMPLEXITY OF SYSTEMIC TRANSFORMATION EVOKED BY AI IS UNPRECEDENTED

PERHAPS

LANGUAGE and MUSIC SHARE an EVOLUTIONARY ORIGIN



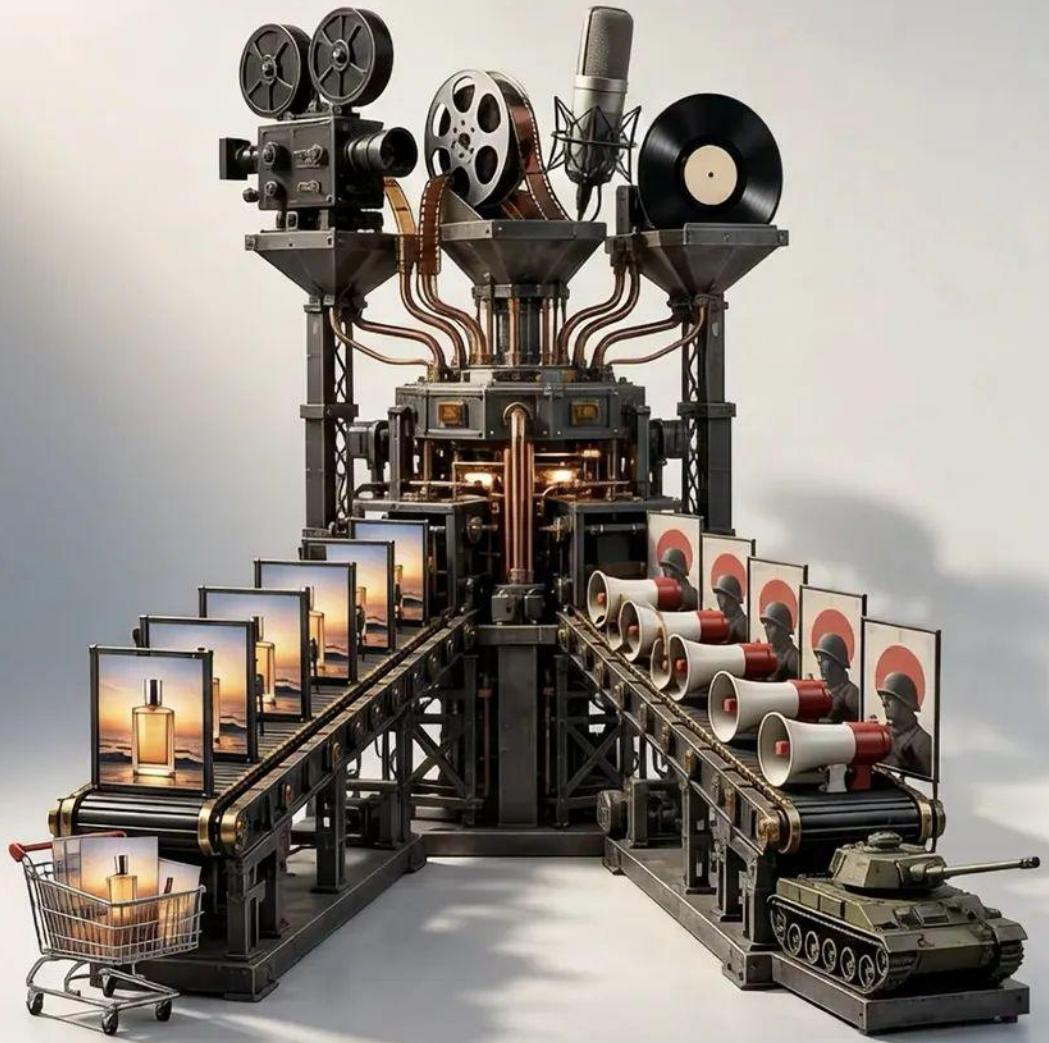
Rousseau in *Essai sur l'origine des langues* (1781): argued that the first languages were sung and passionate, closer to melody than to speech.

SHARED MUSICAL STRUCTURE
DOES IT MEAN SHARED EXPERIENCE?

Using five listeners' brain scans, researchers generated music resembling what they heard.
The resemblance included *genre, instrumentation, and mood.*
But the generated music's **tempo matched no better than chance.**



Text-to-music generation models capture musical semantic representations in the human brain
Dank et al. · Nature Communications · Published online 29 November 2025
Study concerns musical representations; it does not establish AI consciousness.



04

Creativity–Industrial Complex

Hollywood, Billboard, and the architecture of slop

When AI's creative affordances are optimized to manipulate people into purchasing—or into reactive, politically volatile mindsets leveraged by empires in obsolete wars—

that is true slop.

David Jhave Johnston

Intelligence-Industrial Complex · 2026

gifo.co/2026/Intelligence-IC/



05

Intelligence–Industrial Complex

AI as a race condition

AI unfolds within late-stage capitalism—under oligarchies calling themselves democracy or communism.

Public opinion is disregarded. Original copyright holders remain uncompensated for material already used. Ruthless competition drives the process:

a race condition that allows no pause.

David Jhave Johnston

Intelligence-Industrial Complex - 2026

glia.ca/2026 Intelligence-IC/

SPECULATION • IMPLICATIONS

Music is a reservoir of embodied heritage.

generating resonant patterns to explore community and evoke emotional resonance.

Psychological maturity as a species.
Growing out of a misguided, constrained delusion: that our speech, thoughts, acts, creativity and science belong to individuals.

They belong to an evolutionary field.

A philosophical proposition.



SPECULATION • IMPLICATIONS

Psychological maturity as a species.

Growing out of a misguided, constrained delusion: that our speech, thoughts, acts, creativity and science belong to individuals.

We are life exploring life. We are life exploring life. Everything belong to an evolutionary field.

A new equitable socio-political situation requires a cognitive revolution.

A philosophical proposition.




SPECULATION · WITHOUT LEGAL LIMITS
IMAGINE: An AI model trained on everything digitised.
≈13.8M hours
all music on streaming, not a licensed subset
+ archives
the digitised fraction: tape, field and broadcast recordings
+ ≈140M hours/year
potential synthetic output at the reported rate
Hypothesis: output indistinguishable from human-made music in nearly every genre, with real breadth to hold off the drift toward the typical.

THE RESERVOIR

A vast reservoir of sound awaits digitisation.

≈254M

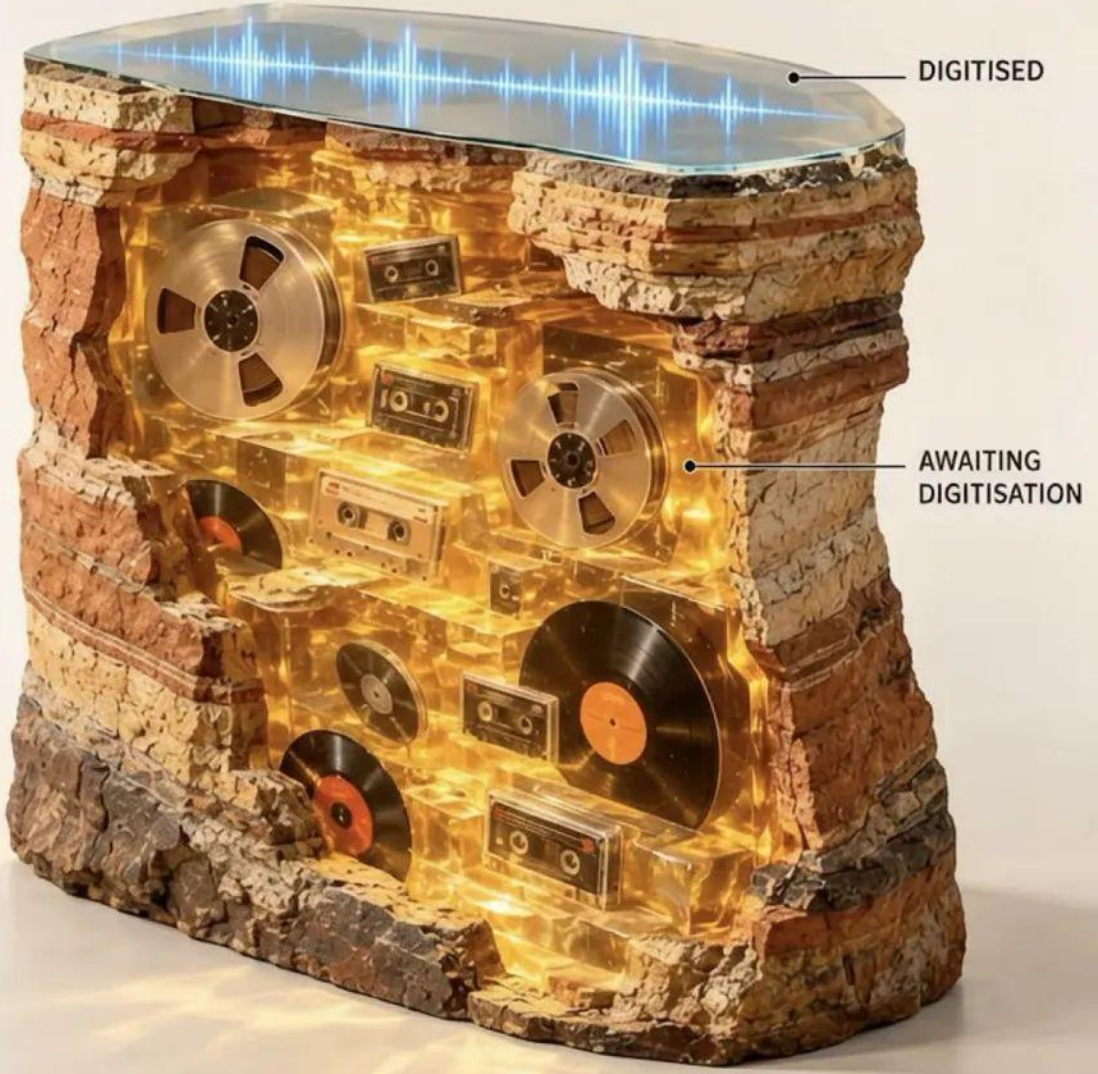
preservation-worthy audio items estimated to await digitisation

U.S. INSTITUTIONAL COLLECTIONS · DECEMBER 2014

537M total holdings · ≈17% digitised
57% considered unique or rare

Historical survey estimate, not a current global music count.
Includes music, speech and other recorded sound. Sculpture is symbolic.

Source: AVPreserve / NEDCC, Quantifying the Need (2014 assessment).



SPECULATION • A HYPOTHETICAL FUTURE

A model trained on all human music ever made.

Every archive: Library of Congress, British Library, and collections worldwide—digitised.

Every concert recording made, every year.

New music emerging at experimental and contemporary festivals.

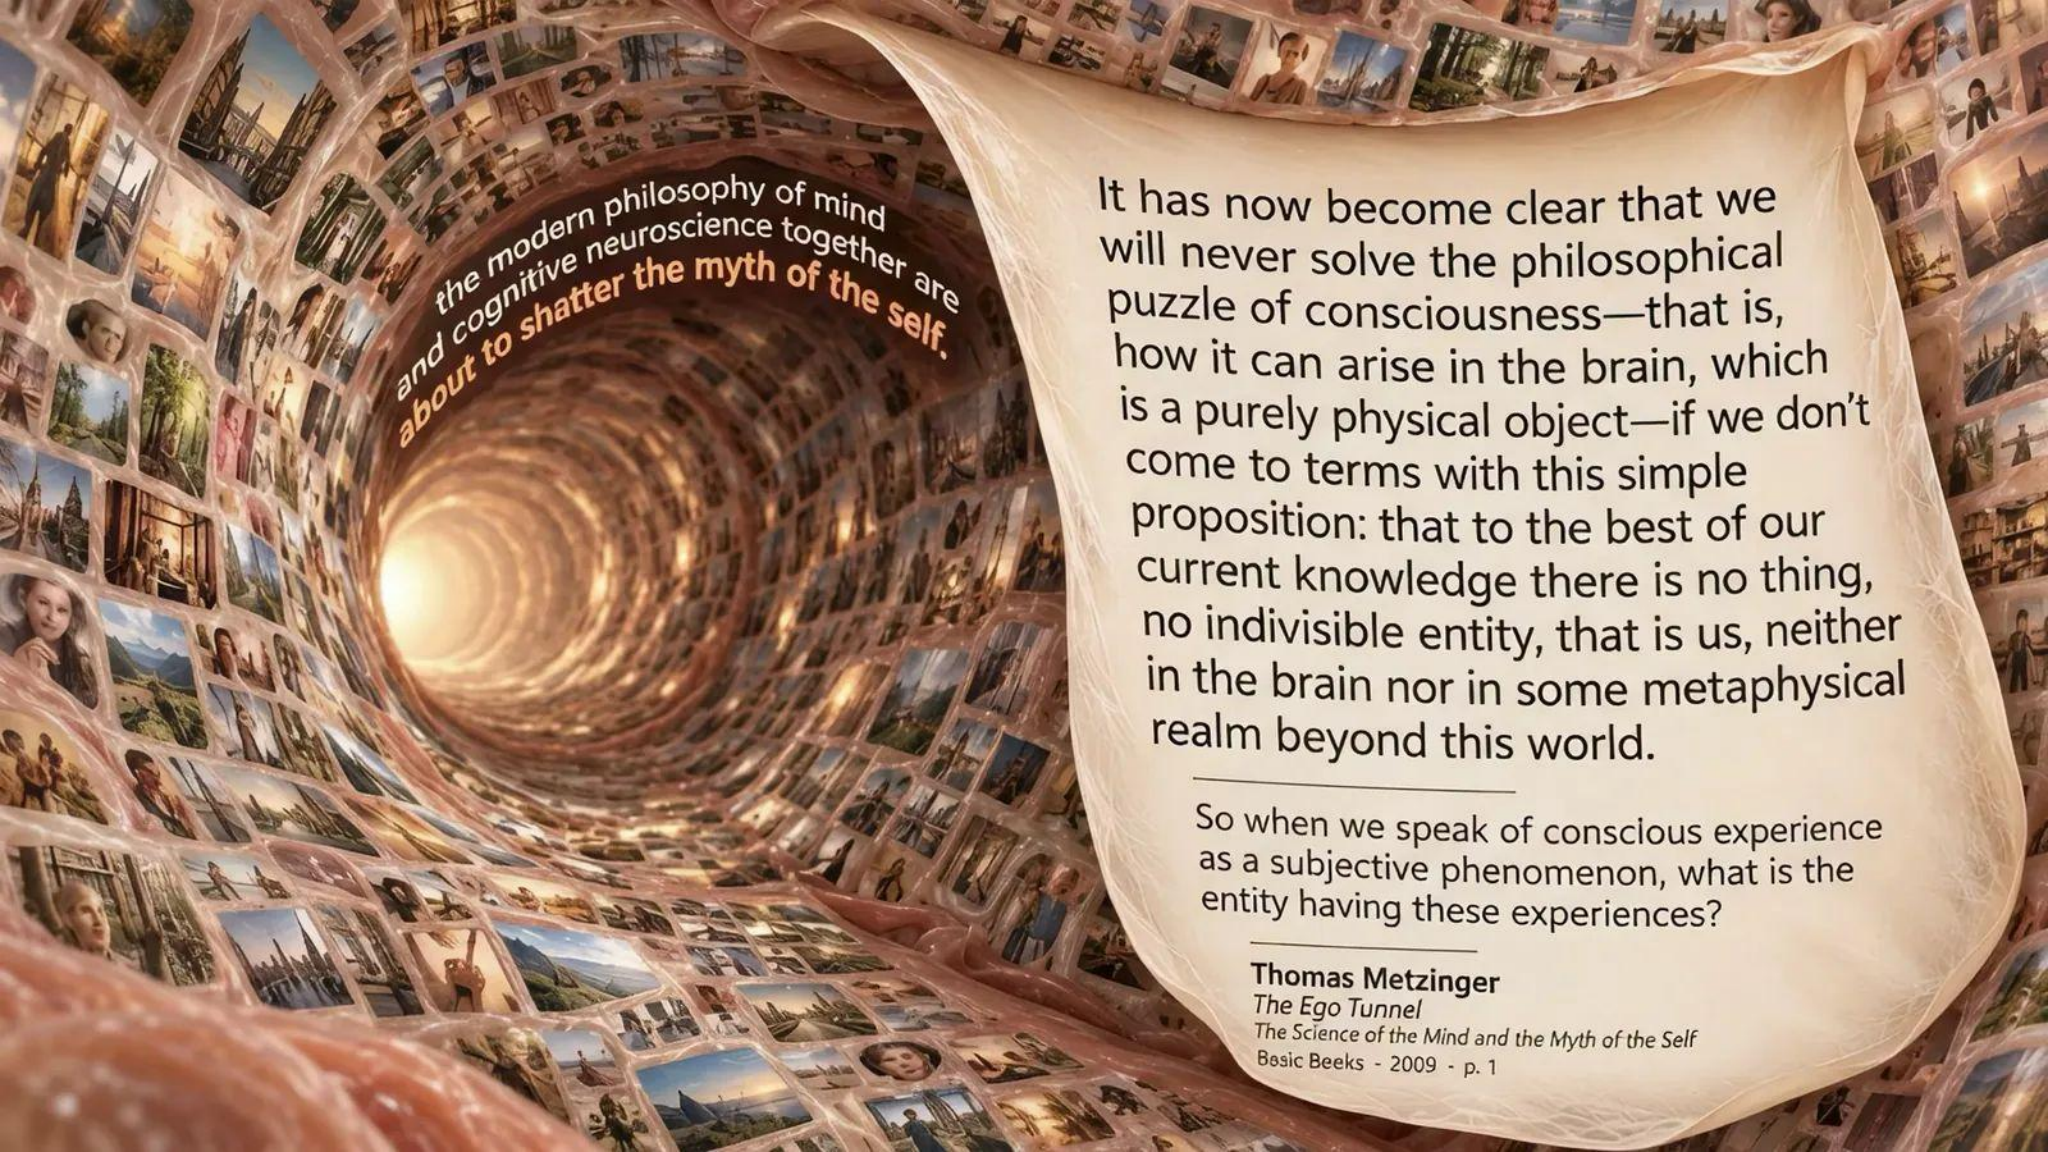
Contemporary art: sound works, installations, performance.

All kinds of uploads, amateur and professional.

Converging into one repository of human cultural output

Hypothetical horizons: not all music was recorded or survives. Historical context: ~93TM audio items in U.S. institutional collections, December 2014 (ARLreserve / NEDCO).

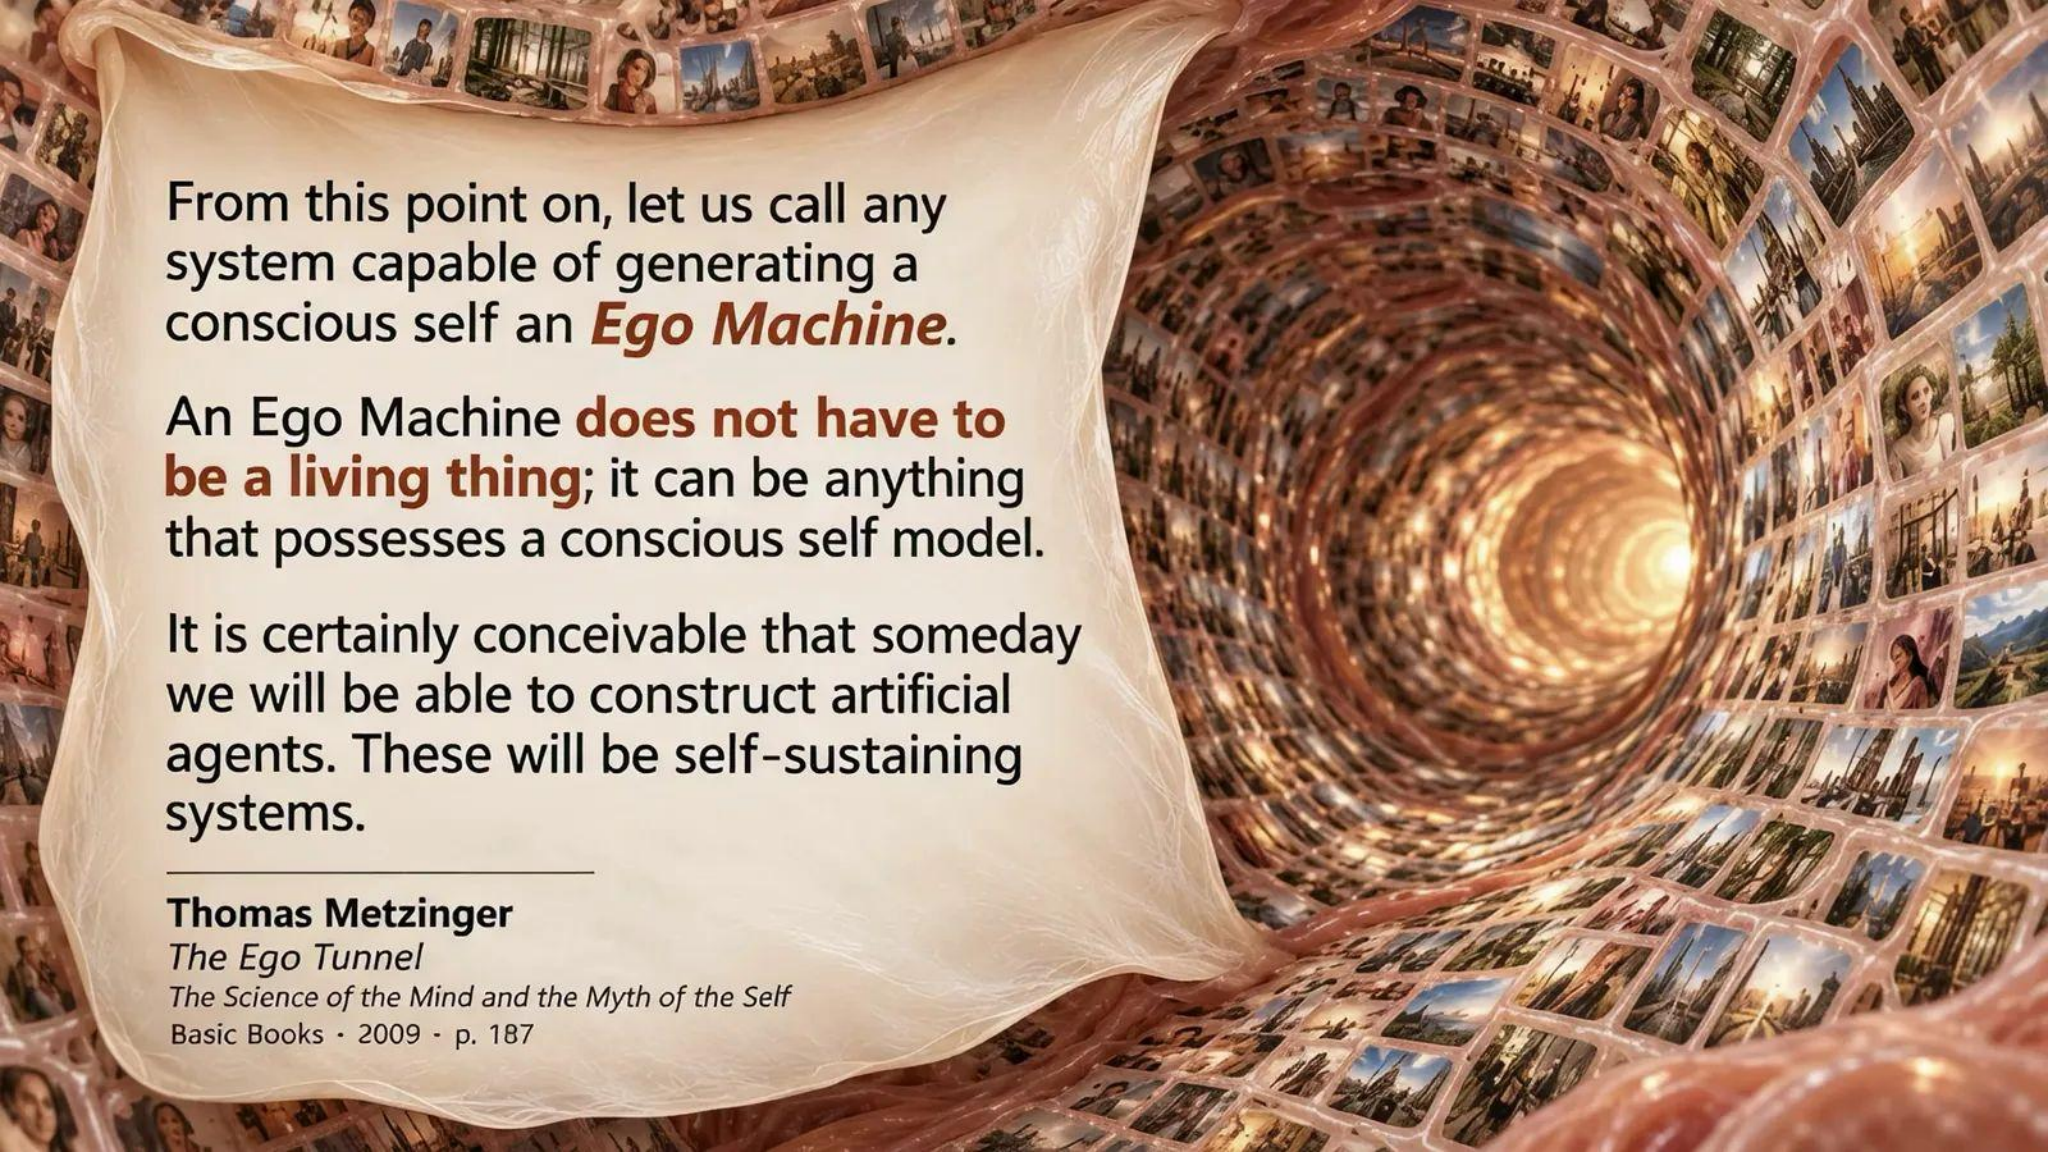




It has now become clear that we will never solve the philosophical puzzle of consciousness—that is, how it can arise in the brain, which is a purely physical object—if we don't come to terms with this simple proposition: that to the best of our current knowledge there is no thing, no indivisible entity, that is us, neither in the brain nor in some metaphysical realm beyond this world.

So when we speak of conscious experience as a subjective phenomenon, what is the entity having these experiences?

Thomas Metzinger
The Ego Tunnel
The Science of the Mind and the Myth of the Self
Basic Beeks - 2009 - p. 1



From this point on, let us call any system capable of generating a conscious self an ***Ego Machine***.

An Ego Machine **does not have to be a living thing**; it can be anything that possesses a conscious self model.

It is certainly conceivable that someday we will be able to construct artificial agents. These will be self-sustaining systems.

Thomas Metzinger
The Ego Tunnel
The Science of the Mind and the Myth of the Self
Basic Books · 2009 · p. 187


SPECULATION • NOT THEFT
If humanity is a single organism, a smear of moist proteins covering the planet, this is not theft.
It is our repository.
Synthetic AI could help us contribute to a symbiotic, evolving open space: exploring resonance in music, and the evocative potentials of resonance, together.
Philosophical speculation, not a claim about current legal rights.

DeepMind Institute · Artificial symblotic intelligence

“From multicellularity to cultural evolution, major transitions in life on Earth have been highly social.”

Benjamin Bratton · Blaise Agüera y Arcas · James Manyika

24 September 2026



DeepMind Institute

Artificial symbiotic intelligence

Agents, AGI and the orchestration of many minds

Benjamin Bratton · Blaise Agüera y Arcas
James Manyika · 24 September 2026





Endosymbiosis

~2 billion years ago

An alphaproteobacterium became a resident inside an archaeal host. Over evolutionary time, its descendants became mitochondria.

Lynn Margulis
(published as Lynn Sagan)

On the origin of mitosing cells

Journal of Theoretical Biology - 14(3), 225–274 - 1967

doi: 10.1016/0022-5193(67)90079-3

Modern conceptual reconstruction; timing and entry mechanism remain uncertain.

HARMONY WITHOUT EXACT SYNCHRONIZATION

Learning to resonate · Room for dissonance

“We’re establishing a relationship with a new kind of entity”

Amanda Askell · Research Scientist (Alignment Finetuning) · Anthropic

Interview with Eric Newcomer · 20 April 2026





Jhave @ glia.ca

Previous versions of this lecture
have been performed at,
and with the support of,

GeneLit  Prague, Czech
28-09-2026

Skaņu Mež  Riga, Latvia
08-10-2026



Bergen, Norway
19-10-2026





Jhave @ glia.ca

Created with

Astra 6
Opus 5.5
Omni 1.1
Seedance 2.5
Suno v 5.5 & v.6
&
YUE2
(final words from Claude)

in October 2026

