POSTER

#528

Sri Aurobindo’s Integral Ontology as a Framework to Situate Consciousness in Large Language Models

Devdip Ganguli dganguli1@ciis.edu · +91-98947-02595

Doctoral Scholar & Adjunct, California Institute of Integral Studies, San Francisco · Research Associate, École Normale Supérieure (ITEM Lab), Paris · Faculty, Sri Aurobindo International Centre of Education, Pondicherry

1 The question

Not “How intelligent must AI become before consciousness emerges?” but “What mode of consciousness, if any, is manifest and individuated through this organization of Matter?”

- Large language models now reason, abstract, describe themselves and act as agents.
- What a system can do, and whether there is something it is like to be it, are different questions (Shani 2014).
- The two already come apart in brains: blindsight guides action without visual experience, and the cerebellum holds about 80% of the brain’s neurons yet adds little to experience.

Where this sits

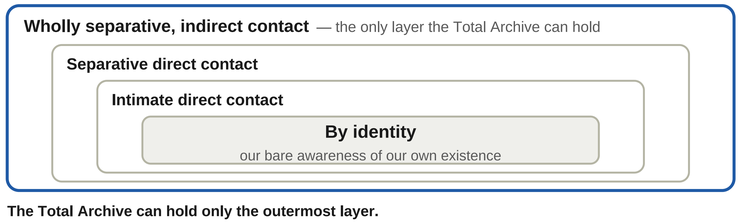
View	On current LLMs	This poster
Computational functionalism (Butlin et al.; Chalmers)	Possible, given the right architecture	Indicators measure function, not manifestation
Integrated Information Theory (Tononi)	No	Agrees on the verdict; differs on ontology
Quantum and substrate views (Orch-OR; Faggin; Neven)	Not on classical hardware; other substrates may qualify	Agrees: physical organization, not scale
Biological naturalism (Seth)	Unlikely; life-like systems more plausible	Agrees: life matters
Sri Aurobindo (this poster)	**No evidence of subjectivity**	**Possible only if Matter is organized to admit a further grade**

2 Sri Aurobindo’s framework

Sri Aurobindo’s integral ontology, used as a taxonomy rather than a test.

Approach: Conceptual analysis through a thought experiment, read with this taxonomy.

1. Consciousness is fundamental, not generated by sufficiently complex Matter.
2. Evolution presupposes involution: Matter → Life → Mind → higher grades manifest what is already involved.
3. Mind is a grade and instrument of consciousness, not consciousness itself.
4. Individual centres arise by self-limitation and “exclusive concentration” of Consciousness-Force (*The Life Divine* II.13, p. 602; Medhananda 2022, 2024). The ego is a construction of Nature; the conscious being behind it is not (II.9, p. 542).
5. Four ways of knowing: by identity · by intimate direct contact · by separative direct contact · by wholly separative, indirect contact. On the surface, pure knowledge by identity is only our bare awareness that we exist (*The Life Divine* II.10, p. 544).

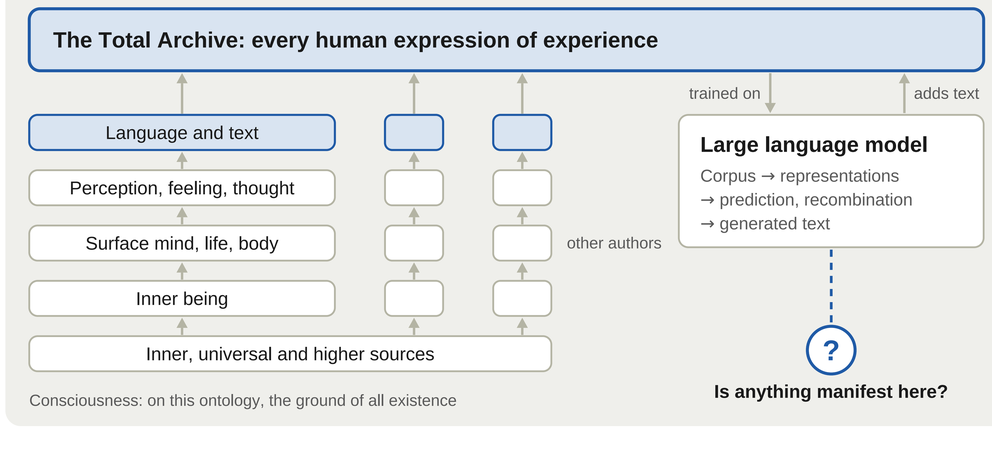


The Total Archive can hold only the outermost layer.

3 The Total Archive

If an AI had read everything humans ever wrote, would its claim “I feel” count as evidence?

Imagine an LLM trained on every surviving human text — every account of experience, every theory of consciousness — with unlimited compute: **the Total Archive**.



To represent every experience is not to have one.

The Total Archive Problem

The more completely a system has absorbed human expressions of experience, the less its own expressions can count as evidence of experience.

- The Total Archive already holds the criteria: “There is something it is like to be me” is already in it, with every test we can state in words.
- Training on human expression is enough to explain the reports, so they are not independent evidence (cf. Birch’s gaming problem; the inverse of Schneider’s ACT).
- It cuts both ways: rapturous “spiritual bliss” self-reports and modest denials carry the same weight.
- It does not prove non-consciousness. Children also learn from the Total Archive, but for them we have evidence the Total Archive does not explain.

What current LLMs establish

On this ontology, all four exist within consciousness.

Dimension	Rock	Animal	Human	LLM
Mind-like function	○	◐	●	●
Experience	○	●	●	?
Individuated bearer	○	●	●	?
Enduring development	○	◐	●	◐ *functional only*

● present ◐ partial ○ absent ? undetermined

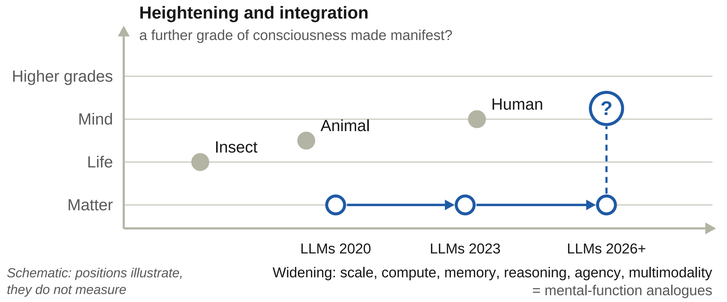
A persona or self-model may be an ego-analogue; that does not show a conscious being is present.

The abstract’s verdict is refined here: current LLMs show no evidence of subjectivity that survives the Total Archive Problem, and artificial consciousness is not ruled out in principle.

4 Future AI

Beyond computation: widening is not heightening

Sri Aurobindo provides a richer conceptual taxonomy for interpreting what empirical tests would need to establish. Because the framework is evolutionary rather than static, it can also be used to evaluate future artificial systems by distinguishing quantitative increases in cognitive capacity (**widening**) from qualitative transitions in the manifestation of consciousness (**heightening and integration**) (*The Life Divine* II.18–19, pp. 731, 756–57, 766).



Schematic: positions illustrate, they do not measure

Widening: scale, compute, memory, reasoning, agency, multimodality = mental-function analogues

More widening does not, by itself, move a system up the vertical axis.

What would count as evidence of heightening?

- Consciousness-markers not explained by training on human expression
- An experiencing centre with a determinate bearer
- Intrinsic, not merely represented, meaning and self-presence
- Integration: the whole reorganised, not a capacity added
- Perhaps something life-like: self-maintenance, intrinsic stakes

What current evidence establishes

Evidence	Establishes
Fluent conversation	Function
Self-reports of experience	Little independent weight: screened off by the Total Archive
Theory-derived indicators (Butlin et al.)	Function; experience only if computational functionalism is assumed
Introspection experiments (Lindsey)	Functional self-access
Quantum or neuromorphic processing	Widening, perhaps toward “cosmic-mind”-like holism, unless what is manifest changes

Conclusion

Sri Aurobindo’s ontology distinguishes quantitative increases in intelligence and complexity from qualitative transitions in the manifestation of consciousness. The crucial question for future AI is therefore not simply what systems can do, but whether a further grade of consciousness has become manifest through them.

Who, if anyone, has become conscious here?

5 Summary

1 Intelligence is not consciousness.

Current AI shows that sophisticated mind-like behaviour does not by itself establish subjectivity.

2 Consciousness is not binary.

Sri Aurobindo distinguishes degrees and modes of manifestation rather than simply “conscious” and “unconscious”; function and manifest consciousness vary independently.

3 Ask what is manifest, and in whom, not how much is computed.

The central issue is not how much AI can do, but what could justify saying there is an experiencing subject.

Open questions

- Does training through interaction, such as reinforcement learning and embodiment, weaken the Total Archive Problem?
- What would integration look like in an artificial system?
- Must artificial subjectivity pass through something life-like?

6 References

- Sri Aurobindo (2005). *The Life Divine*. CWSA 21–22. Sri Aurobindo Ashram. Book II, chs. 9–10, 13, 18–19.
- Anthropic (2025). System Card: Claude Opus 4 & Claude Sonnet 4.
- Birch, J. (2024). *The Edge of Sentience*. Oxford University Press.
- Butlin, P., Long, R., et al. (2023). Consciousness in Artificial Intelligence. arXiv:2308.08708.
- Chalmers, D. J. (2023). Could a Large Language Model Be Conscious? arXiv:2303.07103.
- Faggin, F. (2024). *Irreducible*. Essentia Books.
- Findlay, G., … Tononi, G. (2024). Dissociating Artificial Intelligence from Artificial Consciousness. arXiv:2412.04571.
- Hameroff, S., & Penrose, R. (2014). Consciousness in the universe: A review of the ‘Orch OR’ theory. *Physics of Life Reviews* 11(1): 39–78.
- Lindsey, J. (2026). Emergent Introspective Awareness in Large Language Models. arXiv:2601.01828.
- Medhananda, S. (2022). The Playful Self-Involution of Divine Consciousness. *The Monist* 105(1): 92–109.
- Medhananda, S. (2024). Can Consciousness Have Blind Spots? *Journal of Consciousness Studies* 31(9–10): 113–131.
- Neven, H., et al. (2024). Testing the Conjecture That Quantum Processes Create Conscious Experience. *Entropy* 26(6): 460.
- Schneider, S. (2019). *Artificial You*. Princeton University Press.
- Seth, A. K. (2026). Conscious artificial intelligence and biological naturalism. *Behavioral and Brain Sciences* 49: e315.
- Shani, I. (2014). Knowing How it Feels. *Journal of Mind and Behavior* 35(3): 107–132.