

Article

Not peer-reviewed version

What Artificial Intelligence Is Missing – and Why It Cannot Attain It Under Current Paradigms

[Pavel Straňák](#)*

Posted Date: 14 August 2025

doi: 10.20944/preprints202508.1005.v1

Keywords: artificial intelligence; consciousness, abduction; autopoiesis; phenomenal experience; philosophy of mind; simulation vs. reality; AGI; ontology of intelligence; biological autonomy; meaning-making



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

What Artificial Intelligence Is Missing—And Why It Cannot Attain It Under Current Paradigms

A Philosophical Inquiry into Consciousness, Computability, and Living Systems

Pavel Straňák

Czech Radio, Vinohradská 12, Prague, Czech Republic, EU; pavel.stranak@rozhlas.cz

Abstract

Contemporary artificial intelligence (AI) achieves remarkable results in data processing, text generation, and the simulation of human cognition. However, it fundamentally lacks key characteristics of living systems — consciousness, autonomous motivation, and genuine understanding of the world. This article critically examines the ontological divide between simulated intelligence and lived experience, using the metaphor of the motorcycle and the horse to illustrate the blindness of technological progress. Drawing on philosophical concepts such as abduction, tacit knowledge, phenomenal consciousness, and autopoiesis, the paper argues that current approaches to developing Artificial General Intelligence (AGI) may overlook deeper principles of life and mind. Methodologically, it employs a comparative ontological analysis grounded in philosophy of mind, cognitive science, systems theory, and theoretical biology, supported by contemporary literature on consciousness and biological autonomy. The article calls for a new paradigm that integrates these perspectives — one that asks not only “how to build smarter machines,” but “what intelligence and consciousness truly are.”

Keywords: artificial intelligence; consciousness; abduction; autopoiesis; phenomenal experience; philosophy of mind; simulation vs. reality; AGI; ontology of intelligence; biological autonomy; meaning-making

1. Introduction: The Motorcycle and the Horse—A Metaphor for Overlooked Principles

I arrived at the conference on a motorcycle — fast, efficient, and powerful. It surpassed the horse in every measurable way — except one: the spark of life. A horse is born, grows, learns, fears, rejoices, and wills. A motorcycle requires a factory, fuel, and maintenance. It is unable to build itself, learn autonomously, or make self-originating decisions.

In much the same way, today's artificial intelligence outperforms humans in data processing speed, memory capacity, and text synthesis. However, it fundamentally lacks consciousness, motivation, and autonomy. If we focus solely on performance, there is a risk of overlooking the principles that make human intelligence truly alive.

This metaphor reveals a fundamental distinction between **functional imitation** and **authentic intelligence**. While current AI systems achieve impressive results across many tasks, they lack the essential qualities of living beings: spontaneity, autonomous motivation, and genuine understanding of the world — not merely its simulation. The core philosophical thesis of this article is simple: *simulation is not experience*, and no degree of simulation can fully bridge an ontological gap.

2. Methodology: Beyond Algorithms—Abduction, Tacit Knowledge, and the Limits of Simulation

This section outlines the conceptual framework adopted in the present analysis. It combines philosophical argumentation with comparative ontology, integrating concepts from philosophy of mind, cognitive science, systems theory, and theoretical biology. These are critically evaluated against the constraints and capabilities of current AI architectures, with an emphasis on identifying ontological limitations that cannot be resolved by incremental technical improvements.

In *The Myth of Artificial Intelligence*, Erik J. Larson warns against the illusion of progress [1]. He argues that current AI systems cannot precisely replicate the essence of human intelligence — abduction, the intuitive leap that allows us to form hypotheses without complete data [2]. AI does not know what it does not know — and it cannot ask, because it has no awareness of ignorance. It lacks epistemic self-awareness. Worse still, AI systems' successes may hinder deeper inquiry by creating the false impression that we are close to solving intelligence.

Charles Sanders Peirce, the founder of abductive reasoning, described it as the logic of creativity — the ability to “guess” the best explanation from available data [2]. Michael Polanyi later emphasized that much of human knowledge is **tacit** — unspoken, intuitive, and impossible to fully formalize [3].

Larson's critique strikes at the heart of AI research. Abduction — the capacity to generate plausible hypotheses from incomplete information — requires a creative logical leap that current algorithms are unable to perform. Deduction follows rules, induction finds patterns, but abduction invents new possibilities from silent knowledge. It is this capacity that enables scientific breakthroughs, artistic originality, and everyday decision-making under uncertainty.

Machine learning systems model patterns in data, but they cannot explain them. They do not generate unpredictable hypotheses, and when they do, they cannot assess their relevance. They have no goals — only loss functions to optimize. They simulate logic — but they do not think, feel, or intend. As Daniel Dennett noted, “Computers only do what we understand well enough to teach them” [4]. Even if this statement seems outdated in light of recent generative AI advances, its essence remains valid.

Larson argues that the myth of AI “has deeply harmful consequences when taken seriously, because it undermines science” [1]. The current pursuit of AGI relies on refining transformer architectures, adding agentic systems, chains of thought, and memory modules. These are sophisticated simulations of human behavior — but not its principles. They are motorcycles, not horses: fast, efficient, and powerful — but something essential is missing.

3. Discussion: Consciousness, Autopoiesis, and the Ontological Divide

Consciousness probably does not appear to emerge merely from computational complexity. It is ontologically distinct. Thomas Nagel famously asked, “What is it like to be a bat?” — pointing to the irreducible nature of subjective experience, or *qualia* [5].

David Chalmers framed the “hard problem of consciousness” as the challenge of explaining why subjective experience exists at all [6]. AI systems such as large language models can generate texts that describe emotions, but they do not experience what they describe. As John Searle illustrated in his Chinese Room thought experiment, functional simulation of understanding is not the same as genuine understanding [7]. A computer manipulating Chinese symbols may produce coherent output, but it does not comprehend the language. Today's AI models operate similarly — manipulating numerical representations of tokens to produce human-readable text, but without internal grasp of meaning.

Recent theoretical syntheses reinforce this distinction. Seth and Bayne's survey of contemporary theories of consciousness underscores that no prevailing account treats consciousness as reducible to a mere “feature list” implementable in code; rather, consciousness is seen as an integrated, holistic

property of living systems [8]. This aligns with the view presented here: simulation of conscious behavior is not equivalent to conscious experience.

Table 1 summarizes the key structural and functional differences between living systems and artificial constructs. While machines rely on external design, energy, and control, living organisms generate their own structure, motivation, and adaptation from within. This contrast highlights the ontological gap between simulation and genuine autonomy — a gap that current AI architectures have still failed to bridge.

Table 1. Structural and Functional Differences between living systems and artificial constructs.

Aspect	Living Systems (e.g. Horse)	Artificial Machines (e.g. Motorcycle / AI)
Origin	Self-generated (biological reproduction)	Externally constructed (factory, design)
Information Source	Internal (DNA, cellular processes)	External (blueprints, programming, datasets)
Energy Acquisition	Autonomous (metabolism, environment)	Dependent on external input (fuel, electricity)
Self-replication	Yes (reproduction)	No
Self-regulation	Yes (homeostasis, adaptation)	Limited (predefined feedback loops)
Intentionality	Intrinsic (motivation, goals)	Simulated or externally assigned objectives
Consciousness	Present (subjective experience, qualia)	Absent (no phenomenal awareness)
Development	Evolves and learns organically	Updated via external intervention (AI training, retraining, upgrades)
Ontological Status	Autopoietic (self-creating and sustaining)	Heteropoietic (created and maintained from outside)

Living systems possess the ability to self-replicate, self-regulate, and generate intention. AI, in its current form, lacks internal states, motivation, and autonomous origin. It reacts but does not experience it. It functions cognitively, but it does not live.

This is where the concept of **autopoiesis** becomes decisive. Maturana and Varela originally defined autopoiesis as the self-production of a system that maintains and regenerates its own organization [9]. Moreno and Mossio expanded on this, arguing that biological autonomy entails a *causal closure of constraints*, whereby the organism's regulatory structures are generated and maintained internally [10]. This form of organization is categorically absent in artificial systems.

Recent analyses suggest that “no current AI systems are conscious,” yet also claim that “there are no obvious technical barriers to implementing certain features of consciousness artificially” [11]. However, Kauffman and Roli caution that the living world is not reducible to a theorem: its open-ended creativity cannot be fully captured by formal models [12]. From this perspective, the absence of consciousness in AI is not a temporary technical gap but a consequence of the fundamentally different causal architectures of living and artificial systems.

Philosophical analysis also supports the idea that “consciousness is the initiator of motivation” [13]. Without authentic consciousness, present AI cannot develop true autonomous goals. It can follow programmed objectives or mimic motivated behavior, but it cannot spontaneously generate

its own intentions. Froese and Taguchi highlight that without a mechanism for intrinsic meaning-making, any goal structure in AI remains externally imposed and semantically inert [14].

Consciousness is probably not just another module to be added to an architecture. It is a foundational property of living beings. If AI lacks it, then it is not alive — and may never be, unless built on radically new principles of organization.

4. Unifying Principle: What Connects AI and Machine Production

The motorcycle cannot build itself. It requires a factory, machines, plans, and materials. AI likewise depends on data, infrastructure, and human guidance. Life needs none of these — it replicates, grows, and adapts on its own. This suggests the existence of a principle that current science may not yet understand or may be ignoring.

This parallel between AI and industrial production is not accidental. Both rely fundamentally on external organization and the supply of energy and information. Living systems, by contrast, are *autopoietic* — they create and sustain themselves through their own internal processes [9,10]. This capacity for self-organization and self-creation is radically different from the *heteropoietic* nature of current AI. Figure 1 illustrates the key differences between living and artificial systems.

Humberto Maturana and Francisco Varela described living systems as entities whose organization exceeds the sum of their components [9]. They are not merely systems that perform tasks — they are systems that *become*, evolving from within. Present AI systems, by contrast, originate entirely from external processes; they never arise from themselves.



Figure 1. Graphical illustration of the differences between living and artificial systems. A horse "replicates" itself, draws the necessary information "from within," and obtains the required energy and material from its surrounding environment. In contrast, a motorcycle requires blueprints, a factory, machines, and external sources of energy and materials for its production.

Kauffman and Roli emphasize that living systems are characterized by *endless adjacent possible states*, meaning they can continually generate novel functions and structures beyond any pre-

specified state space [12]. This property resists algorithmic prediction and simulation, undermining the assumption that scaling computational complexity will inevitably yield consciousness.

It is possible that the missing principle is internal organization. It might be something akin to “biological intentionality,” or it might be consciousness itself [8,12]. Until we understand this principle, we will continue building better motorcycles — hoping one day they will run and feel like horses.

The prevailing reductionist approach to AI presumes that real intelligence will emerge from increasing complexity and computational power [15]. Yet this view overlooks the possibility that consciousness — and true intelligence — may require qualitatively distinct principles of matter and information organization [8,12,14]. Today’s AI systems are built upon vast arrays of silicon-based logic structures, trained on immense corpora of multimodal data. The blinking LEDs of data centers may evoke the illusion of thought — but in principle, the same function could be replicated by an astronomical number of gears and levers. Such a (practically unbuildable) mechanical apparatus could converse with us just like today’s chatbots. But could a sufficiently vast system of gears actually think? Could it be conscious? And if so — how many gears would be enough? Are we perhaps missing something fundamental? **Figure 2** highlights the illusion of life often subjectively attributed to electronic AI systems.




Figure 2. AI implemented in a data center creates the illusion of life—as if machines had begun to think. Yet in principle, the same results could be achieved using billions of gears performing exactly the same function. In that case, however, the subjective feeling of a “living machine” does not arise.

5. Implications for the Future of AI

This analysis does not imply that current AI is not useful. Motorcycles are useful even though they are not horses; in fact, they fulfill certain functions significantly better. Present AI can solve many practical problems without being conscious or alive — even at levels inaccessible to humans. A simple calculator already surpasses human capacity in numerical precision and speed. However, if our goal is to create truly intelligent partners, we must reconsider our fundamental assumptions.

Large Language Model (LLM) AI systems simulate human conversation with artificial yet highly convincing fluency. They create an illusion of understanding that does not actually exist, amplified by our innate tendency to anthropomorphize anything that behaves in a human-like manner.

As Luciano Floridi observed, “AI is not about machines being like us. It’s about how the world changes when parts of reality begin to act like us” [16].

Recent debates about whether systems like ChatGPT are conscious highlight the urgency of this distinction. If we cannot separate sophisticated simulation from authentic consciousness, we risk both overestimating current systems and underestimating the complexity of true intelligence. Table 2 summarizes the conceptual limitations of current AI systems.

Table 2. Conceptual Limits of AI: A Comparative Overview.

Concept	Definition	Relevance to Living Systems	Limitation in AI Systems
Autopoiesis	Self-creation and self-maintenance of a system through internal processes	Living organisms generate and sustain themselves autonomously	AI systems are externally designed and lack self-generative structure
Abduction	Reasoning that generates hypotheses from incomplete data	Humans intuitively infer meaning and possibilities beyond given inputs	AI lacks genuine hypothesis formation; it extrapolates from training data
Phenomenal Consciousness	Subjective experience – the “what it is like” aspect of being conscious	Living beings possess first-person awareness and qualia	AI has no subjective experience or inner perspective

The real risk is not that AI will become conscious, but that we will cease to distinguish between simulation and experience.

This has **ethical consequences**: we might attribute moral status, rights, or decision-making authority to entities that are in fact unconscious, while simultaneously underestimating the complexity and fragility of actual conscious life forms. Froese and Taguchi [14] caution that such misattribution could distort the way we value meaning-making processes, both in humans and in artificial agents.

It is possible that we need a radically different approach — not merely scaling current architectures, but making a qualitative leap toward new principles. Perhaps we must first understand how life emerges from inorganic matter before we can create truly living intelligence.

6. Conclusion: Simulation Is Not Experience

Despite the remarkable progress in artificial intelligence, a fundamental ontological gap remains between simulated cognition and lived experience. The metaphor of the motorcycle and the horse illustrates this divide: while machines can outperform biological organisms in speed and efficiency, they lack the intrinsic qualities of life — self-generation, autonomous motivation, and phenomenal consciousness.

This article has argued that current AI paradigms, grounded in external design and optimization, cannot replicate the internal organization and meaning-making processes characteristic of living systems. The absence of autopoiesis and subjective experience in artificial constructs suggests that intelligence, as we understand it in biological terms, may not be attainable through mere computational scaling.

Contemporary AI is immensely useful, but its successes should not lead us to conclude that the path to AGI has already been found. It is possible that genuine intelligence requires principles we do

not yet understand — perhaps ones that touch upon the very nature of consciousness. Whether such principles exist, and whether they can be technically realized, remains an open question.

If consciousness is not an emergent property of complexity alone, but a qualitatively distinct phenomenon rooted in biological organization, then the pursuit of AGI under current paradigms may be fundamentally limited. The challenge is not simply to simulate thought, but to understand what thought and consciousness truly are — and whether they can arise outside the domain of life.

7. Afterthought: Rethinking the Foundations

Recent advances in generative AI have reignited debates about the nature of intelligence and the boundaries of simulation. While these systems exhibit impressive capabilities, they remain constrained by the entropy of training data and the absence of intrinsic intentionality. Theoretical insights from information theory and systems biology suggest that this limitation is not merely technical but ontological, as discussed for example in [17].

Conversely, some theories argue that consciousness can emerge from complexity [18], which clearly indicates that such foundational questions remain open. This paper contributes to the ongoing discourse by offering the possibility that consciousness and autonomy may require principles beyond computation — principles rooted in self-organizing, living systems. While some argue that consciousness could eventually be simulated through sufficient complexity, such claims remain speculative. The fact that 2 kilograms of elementary particles, when arranged in a particular configuration (brain), give rise to conscious experience suggests that material organization matters — but whether computable artificial systems can replicate this remains an open question. The horse does not need to be fast to be alive; the motorcycle does not need to be alive to be fast. Perhaps the real question is not how to simulate intelligence and consciousness, but whether we truly understand what intelligence and consciousness are.

The key proposition of this paper is the possibility that there are fundamental differences between machines and living systems that cannot be adequately described within the current paradigm.

Acknowledgments: The author thanks his colleagues for their valuable insights and discussions that helped shape the ideas presented in this paper. This research was conducted independently, without institutional funding. Some passages of this manuscript — including the figures — were prepared or refined with the assistance of a large language model (LLM). The author takes full responsibility for the content and conclusions presented herein.

Conflict of Interest: The author declares that no competing interests exist.

Data Availability Statement: Not applicable. This manuscript presents a theoretical framework and does not report empirical data.

Funding: This research received no external funding. The work was conducted independently by the author, who is employed by Czech Radio, but the research was carried out privately and outside of institutional duties.

References

1. Larson, E. J. (2021). *The Myth of Artificial Intelligence*. Harvard University Press.
2. Peirce, C. S. (1931–1958). *Collected Papers of Charles Sanders Peirce*. Harvard University Press.
3. Polanyi, M. (1966). *The Tacit Dimension*. University of Chicago Press.
4. Dennett, D. C. (1991). *Consciousness Explained*. Little, Brown and Co.
5. Nagel, T. (1974). "What is it like to be a bat?" *The Philosophical Review*, 83(4), 435–450.
6. Chalmers, D. (1995). "Facing up to the problem of consciousness." *Journal of Consciousness Studies*, 2(3), 200–219.
7. Searle, J. R. (1980). "Minds, brains, and programs." *Behavioral and Brain Sciences*, 3(3), 417–44.

8. Seth, A. K., & Bayne, T. (2022). Theories of consciousness. *Nature Reviews Neuroscience*, 23(5), 307–321. <https://doi.org/10.1038/s41583-022-00587-4>
9. Maturana, H. R., & Varela, F. J. (1980). *Autopoiesis and Cognition: The Realization of the Living*. Springer.
10. Moreno, A., & Mossio, M. (2015). *Biological Autonomy: A Philosophical and Theoretical Enquiry*. Springer.
11. Butlin, P., et al. (2023). Consciousness in AI: Insights from Science of Consciousness. arXiv. <https://arxiv.org/abs/2308.08708>
12. Kauffman, S., & Roli, A. (2021). The World is Not a Theorem. *Entropy*, 23(6), 808. <https://doi.org/10.3390/e23060808>
13. Su, J. (2024). Consciousness in AI: A Philosophical Perspective Through Motivation. *Critical Debates in Humanities*.
14. Froese, T., & Taguchi, S. (2019). The problem of meaning in AI. *Minds and Machines*, 29, 263–288. <https://doi.org/10.1007/s11023-019-09495-2>
15. Baars, B. J. (1997). *In the Theater of Consciousness*. Oxford University Press.
16. Floridi, L. (2014). *The Fourth Revolution: How the Infosphere is Reshaping Human Reality*. Oxford University Press.
17. Straňák, P. (2025). Lossy Loops: Shannon's DPI and Information Decay in Generative Model Training. *Preprints.org*. <https://doi.org/10.20944/preprints202507.2260.v1>
18. Feinberg & Mallatt (2020). *Phenomenal Consciousness and Emergence: Eliminating the Explanatory Gap*. *Frontiers in Psychology*. <https://doi.org/10.3389/fpsyg.2020.01041>

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.