

Article

Not peer-reviewed version

The Dual Structure of Mind and Life Can Explain Consciousness and Evolutionary Creativity Beyond Simulation

Pavel Straňák *

Posted Date: 18 November 2025

doi: 10.20944/preprints202511.1272.v1

Keywords: complex systems; adaptive systems; autopoiesis; developmental canalization; Shannon entropy; computational complexity; AI autonomy; consciousness hypothesis; evolutionary directionality; falsifiable model



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

The Dual Structure of Mind and Life Can Explain Consciousness and Evolutionary Creativity Beyond Simulation

Pavel Straňák

Independent Researcher, Prague, Czech Republic; science.stranak@gmail.com

Abstract

We propose that complex systems in mind and life follow a dual architecture: a *computational component* generating structures (e.g., thoughts, biochemical states) and a *directional component* ensuring long-term coherence and creativity. Large language models (LLMs) simulate the former but lack the latter, resulting in entropic decay during autonomous operation—a process governed by Shannon's Data Processing Inequality. Drawing on autopoiesis theory*and developmental canalization, this hypothesis posits that consciousness and evolutionary directionality are ontologically distinct control layers, falsifiable through long-term AI autonomy tests and evolutionary simulations. The framework bridges empirical AI failures with theoretical principles of adaptive complexity, offering testable predictions for complex systems research.

Keywords: complex systems; adaptive systems; autopoiesis; developmental canalization; Shannon entropy; computational complexity; AI autonomy; consciousness hypothesis; evolutionary directionality; falsifiable model

1. Introduction: What the Existence of LLMs Reveals About Thought and Consciousness

Methodological Note: This study employs analogical modeling between artificial and biological systems, grounded in empirical evidence from AI (e.g., entropic decay in autonomous LLM chains; Huang et al., 2024) and evolutionary biology (e.g., developmental canalization; Moreno & Mossio, 2015). While the dual-structure hypothesis includes ontological interpretation, it generates *falsifiable predictions* testable via long-term AI autonomy and evolutionary simulations.

The emergence of large language models (LLMs) demonstrates that human-like text generation and reasoning patterns can be reconstructed *without consciousness* (Searle, 1980). This empirical observation has stimulated interdisciplinary research into the functional boundaries of artificial systems (Butlin et al., 2023). We propose that both mind and life operate via a dual architecture: a computational component generating structures (thoughts, biochemical states) and a directional component enabling evaluation, stability, and creativity. Current LLMs replicate the former but lack the latter, leading to informational noise under prolonged autonomous operation (Huang et al., 2024). Through analogies between cognitive and biological systems, this work introduces a *falsifiable hypothesis*: consciousness and evolutionary creativity are not emergent byproducts but ontologically distinct layers of complex adaptive systems—manifest in life, cognition, and historical dynamics.

An LLM is like a colossal mill that has ground the entire internet. It contains hundreds of billions of “gears” — parameters tuned to linguistic patterns, knowledge structures, and cognitive procedures embedded in human texts. This mill can generate new sentences, answer questions, and propose hypotheses. And yet it does not think — it resembles a giant mechanical cash register with a crank, spinning its gears to produce text. Just gears and levers, realized in silicon chips — nothing more.

This ability to simulate thought without consciousness is key. It suggests that the human mind likely consists of two functional components: one computational, the other conscious (Dennett, 1991).

LLMs mirror the former. But without the latter — without consciousness — the machine becomes, under prolonged autonomous operation, a powerful generator of noise, unable to fully assess the relevance of its own outputs (Larson, 2021).

2. Thought Versus Consciousness: Two Components of the Human Mind

In ordinary situations, the human brain activates an internal “transformer” — a generator of linguistic structures that quietly speaks to itself. This component is computational, predictive, and to some extent algorithmic, functionally akin to the transformer architecture of AI systems.

Running in parallel, however, is consciousness — a component that evaluates, directs, and feels. Consciousness is not merely a passive observer, but an active “canalizer” of thought. It creates bias, motivation, and direction (Baars, 1997). Without it, human thinking would resemble Brownian motion of words, lacking long-term orientation.

AI has advantages: it accesses all available knowledge, never forgets, and generates rapidly. But it lacks direction. It has no inner drive. It is like a motorcycle without a rider — fast, but without knowing where to go. A human walks, but knows the destination. AI sees the entire “playing field” of human knowledge brightly lit and races across it. A human sees only the small portion they have studied, walks slowly, and lights the path with a flashlight. They move slowly, but they know where — and that is what matters (Floridi, 2014).

3. Evolution as a Dual Process: Chemistry + Canalization

The evolution of life is often described as the result of a blind process — a combination of mutation, selection, and time. Yet this description is extremely improbable if taken literally (Kauffman & Roli, 2021). The likelihood that complex living systems would arise purely by chance is, according to some estimates, so low that it cannot be considered realistic.

This suggests that evolution involves a form of “canalization” — a directional tendency first formalized by Waddington (1942) and later developed in autopoietic frameworks (Moreno & Mossio, 2015). Chemistry and genetics form the strong component of the system — the part we know and measure. But there may also be a second component — weaker, subtler, yet essential (Maturana & Varela, 1980).

When we ignite paper, it burns according to chemical laws. These laws are strong, dominant, and determine the outcome. But in living systems, something else occurs: a structure emerges that repairs itself, replicates itself, and creates itself. Life does not require a factory. All its parts are “alive”; even the periphery regenerates according to an invisible “internal blueprint.” That is why life remains stable over time and continues to evolve. Machines, by contrast, require maintenance after a while and cannot operate independently for long without failure.

4. Analogy: Mind and Life as Dual-Component Systems

A comparison of the principal differences between organisms and artificial systems is presented in Table 1.

Table 1. Functional comparison between artificial systems and living organisms. This table outlines the structural and functional differences between machines and living systems. Both exhibit components that generate and regulate behavior, yet only living organisms appear to possess an internal directional principle – consciousness or developmental canalization – absent in artificial systems. While machines rely solely on computational, mechanical, and chemical operations, living systems may embody additional formative dynamics that remain scientifically undescribed.

Property/Component	AI System/Machine	Living Organism/Mind	Empirical Evidence
Computational Component	Internal LLM, control algorithms, predictive structures	Genetics, biochemistry, cellular processes	AI: Transformer architecture (Vaswani et al., 2017); Life: Central dogma of molecular biology (Crick, 1970)
Directional Component	Absent (no consciousness or intrinsic motivation)	Consciousness, subjective experience, internal canalization	AI: (Huang et al. 2024" – noise in autonomous LLM chains; Life: (Kauffman & Roli 2021) – non-random evolutionary order
Function	Generates outputs via loss function optimization	Generates structures, behavior, creativity	AI: (Larson 2021) – competence without comprehension; Life: "Maturana & Varela 1980" – autopoietic self-creation
Stability	Requires maintenance, degrades, lacks autoregulatory capacity	Self-repair, adaptation, evolutionary continuity	AI: (Straňák 2025) – Shannon's DPI in generative loops; Life: (Moreno & Mossio 2015) – biological autonomy
Ontological Status	Heteropoietic (constructed externally)	Autopoietic (self-constituting within)	AI: (Searle 1980) – syntactic simulation; Life: (Maturana & Varela 1980) – operational closure

If this analogy holds, then consciousness and "evolutionary direction" are not mere side effects of computation or chemistry, but distinct components of reality (Polanyi, 1966).

5. A Critique of Reductionism: Why "More Computation" Is Likely Not Enough

Reductionism is the belief that complex phenomena can be fully explained as the sum and interaction of simpler parts. In the domain of AI, it is often assumed that consciousness, motivation, or creativity will "emerge" if the computational system becomes sufficiently complex, or employs sophisticated techniques such as chains of reasoning (Larson, 2021).

This assumption, however, encounters a fundamental problem: a computational component without direction produces noise. In communication channels – for instance, in digital signal processing – it is well established that no signal processing can increase the amount of information contained in the signal. This is known as Shannon's Data Processing Inequality (Shannon, 1948; Straňák, 2025). Signals tend toward entropy. In AI, the situation is analogous: without consciousness, without an evaluative mechanism, without intrinsic motivation, the outputs of extended reasoning become Brownian motion of words – random noise. This phenomenon has been empirically observed in AI systems many times (Huang et al., 2024).

Contemporary LLM-based AI systems are powerful, but they lack an inner goal. They optimize a loss function, but do not know why. They cannot assess the relevance of their outputs. They cannot effectively self-correct or reorient when they err (Froese & Taguchi, 2019).

Machines wear down, require servicing, and their outputs gradually dissolve into noise. Life, by contrast, creates, repairs itself, and evolves. This suggests that living systems contain something beyond the computational component — something that current science has yet to describe, but which appears essential for stability and creativity. A machine with a sophisticated control computer is, in essence, only half of life: the computational and mechanically predictive half. The other half — consciousness and the canalization of structural stability and development — is missing (Maturana & Varela, 1980), see Figure 1.

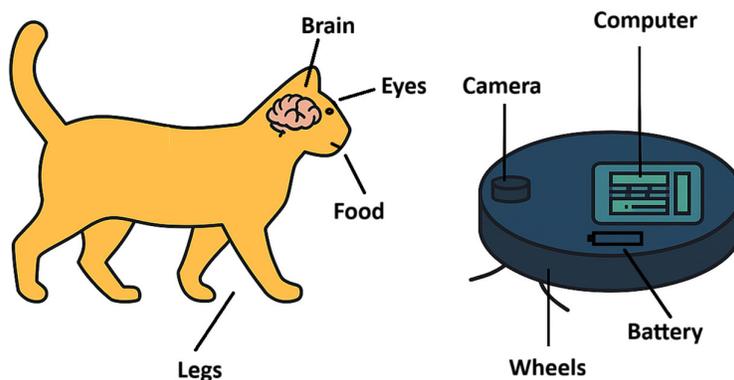


Figure 1. Comparison of corresponding components in life and machines — here, a cat and a robotic vacuum cleaner. Both contain functionally analogous parts; however, in machines these rely solely on computational, mechanical, and chemical principles. In living systems, the directional component may manifest as measurable deviations from purely computational predictions — a hypothesis testable via long-term AI autonomy testing or evolutionary simulations.

6. A Cautious Hypothesis: There May Exist a Formative Principle We Do Not Yet Know

Consciousness and evolutionary direction may not be mere side effects of computational complexity. They may not be computational in nature at all. It is possible that they represent a different kind of organization of reality — one that current science is not yet able to describe, or perhaps systematically overlooks because it does not fit the prevailing paradigm (Feinberg & Mallatt, 2020).

Science is grounded in what can be measured, modeled, and simulated. But what if there are phenomena that elude this methodology? What if consciousness is not “something that emerges,” but something ontologically distinct — much like the difference between a map and the landscape it depicts (Chalmers, 1995)? Or perhaps such phenomena are measurable, but their intensity is so faint that they continue to be overlooked. Especially when they fall outside the paradigm, and thus few are actively searching for them.

In the case of life, it appears that chemistry alone may not suffice. Genes may not be merely a blueprint, but primarily a form of evolutionary memory. Evolution may not be a random walk through a space of possibilities, but rather a process with an internal direction that we do not yet understand (Moreno & Mossio, 2015).

In the case of mind, it appears that computation is likely not enough. An AI LLM can efficiently generate relevant texts, but it does not know what it is saying. It lacks motivation, direction, and experience. This becomes especially evident during prolonged autonomous operation, when the likelihood of losing coherence, hallucinating, and producing informational noise increases (Huang et al., 2024).

We see its manifestations: in life, in mind, in history. This view resonates with Peirce’s conception of reality as inherently semiotic and generative — a layered process in which meaning, direction, and novelty emerge through structured interaction rather than mere computation (Peirce, 1931–1958).

7. Formalization of the Hypothesis Core

To formalize the dual-structure hypothesis, we propose a minimal dynamical model distinguishing a *computational component* that generates structures and a *directional component* that evaluates and stabilizes long-term coherence. Let $x(t)$ represent the system state at time t (e.g., token sequence in an LLM, biochemical configuration in a cell), evolving via a known generative function f with parameters θ :

$$x(t+1) = f(x(t), \theta + \beta \cdot d(t)) \quad (1)$$

The directional state $d(t)$ encodes an internal bias toward relevance or fitness, updated slowly via a relevance gradient $\nabla R(x(t))$:

$$d(t+1) = d(t) + \alpha \cdot \nabla R(x(t)) \quad (2)$$

where $\alpha \ll 1$ and $\beta > 0$ in living or conscious systems, but $\beta = 0$ in current LLMs. When $\beta = 0$, the dynamics reduce to pure computation, and Shannon's Data Processing Inequality implies that mutual information with any distal goal monotonically decreases—predicting the observed entropic collapse in autonomous LLM chains (Huang et al., 2024). In contrast, $\beta > 0$ enables sustained coherence and creative exploration, consistent with biological canalization and conscious deliberation. This model is falsifiable: long-term autonomy tests should reveal whether coherence degrades exponentially without a learnable $d(t)$, offering an empirical bridge between AI limitations and the ontology of adaptive complexity, see also Figure 2.

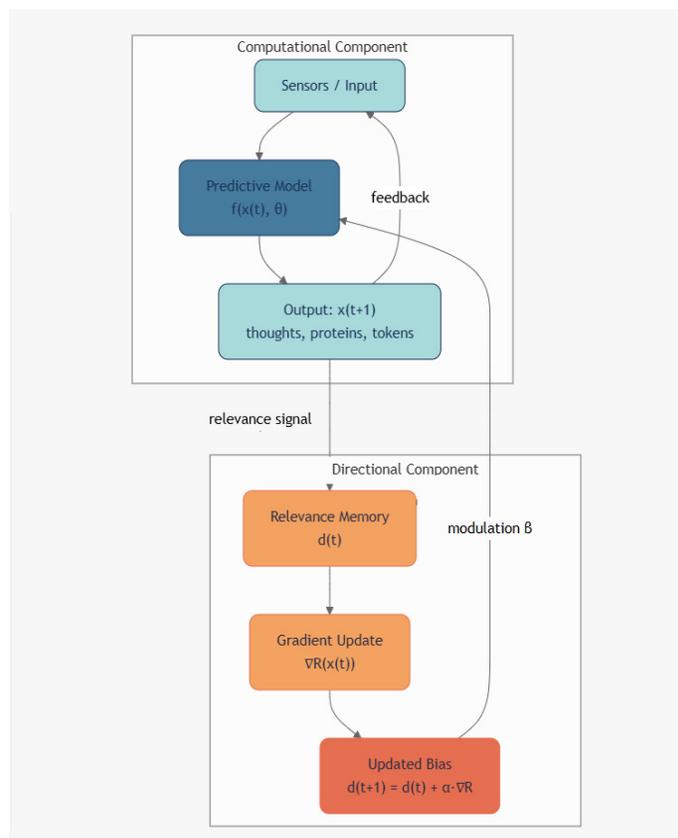


Figure 2. Network diagram of the dual-structure hypothesis in complex adaptive systems. Nodes represent functional layers: the *computational component* (blue) generates structures via predictive transformations (e.g., next-token prediction in LLMs, biochemical synthesis in cells), while the *directional component* (red) evaluates relevance, maintains

coherence, and enables long-term stability via slow adaptation of an internal bias $d(t)$. Solid arrows indicate forward generation and feedback within layers; dashed arrows denote cross-layer modulation with strength β . In artificial systems ($\beta=0$), the loop is open, leading to entropic decay under autonomy (Shannon's DPI). In living or conscious systems ($\beta>0$), closed-loop canalization supports autopoiesis, self-repair, and creative exploration. This architecture formalizes the dynamical model in Equations 1–2 and predicts measurable deviations from purely computational trajectories in long-term behavioral or evolutionary simulations.

8. Conclusion and Research Implications

LLM-based systems empirically demonstrate that linguistic and cognitive patterns can be generated without consciousness. This work proposes a dual-structure hypothesis for complex adaptive systems: a computational layer (replicated in AI) and a directional layer (absent in current models), responsible for long-term coherence and creativity.

The formal model (Equations 1–2) predicts that in the absence of directional modulation ($\beta=0$), autonomous systems undergo entropic collapse, consistent with observed degradation in LLM chains (Huang et al., 2024) and Shannon's Data Processing Inequality. Conversely, $\beta>0$ supports canalized, autopoietic dynamics, enabling sustained coherence and adaptive creativity—features characteristic of biological evolution and conscious cognition.

Research implications: Future empirical studies should evaluate long-term autonomous performance in AI systems (e.g., ≥ 1000 -step reasoning chains without external feedback) and investigate non-random canalization patterns in evolutionary simulations. Exponential decay in output coherence in the absence of a learnable directional state $d(t)$ would provide support for the dual-structure hypothesis; conversely, sustained coherence without such a mechanism would falsify it. This framework thus establishes a rigorous, falsifiable link between observed limitations in artificial autonomy and the theoretical foundations of adaptive complexity in biological and cognitive systems.

Funding: This research received no external funding.

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

Acknowledgments: The author thanks colleagues for discussions that shaped this work. This research was conducted independently without institutional funding. Some passages of this manuscript 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.

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

References

- Baars, B. J. (1997). *In the theater of consciousness: The workspace of the mind*. Oxford University Press.
- Butlin, P., Chella, A., Cheung, B., Crosby, M., Dehaene, S., Halina, M., ... & Yudkowsky, E. (2023). Consciousness in artificial intelligence: Insights from the science of consciousness. *arXiv*. <https://arxiv.org/abs/2308.08708>
- Chalmers, D. J. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies*, 2(3), 200–219.
- Crick, F. (1970). Central dogma of molecular biology. *Nature*, 227(5258), 561–563. <https://doi.org/10.1038/227561a0>
- Dennett, D. C. (1991). *Consciousness explained*. Little, Brown and Company.
- Feinberg, T. E., & Mallatt, J. M. (2020). Phenomenal consciousness and emergence: Eliminating the explanatory gap. *Frontiers in Psychology*, 11, 1041. <https://doi.org/10.3389/fpsyg.2020.01041>
- Floridi, L. (2014). *The fourth revolution: How the infosphere is reshaping human reality*. Oxford University Press.
- Froese, T., & Taguchi, S. (2019). The problem of meaning in AI: Formal models and biological behavior. *Minds and Machines*, 29(3), 385–408. <https://doi.org/10.1007/s11023-019-09501-7>

- Huang, L., Yu, W., Ma, W., Zhong, W., Feng, Z., Wang, H., ... & Liu, T. (2024). A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *arXiv*. <https://arxiv.org/abs/2311.05232>
- Kauffman, S., & Roli, A. (2021). The world is not a theorem. *Entropy*, 23(6), 808. <https://doi.org/10.3390/e23060808>
- Larson, E. J. (2021). *The myth of artificial intelligence: Why computers can't think the way we do*. Harvard University Press.
- Maturana, H. R., & Varela, F. J. (1980). *Autopoiesis and cognition: The realization of the living*. D. Reidel Publishing Company. <https://doi.org/10.1007/978-94-009-8947-4>
- Moreno, A., & Mossio, M. (2015). *Biological autonomy: A philosophical and theoretical enquiry*. Springer. <https://doi.org/10.1007/978-94-017-9837-2>
- Nagel, T. (1974). What is it like to be a bat? *The Philosophical Review*, 83(4), 435–450. <https://doi.org/10.2307/2183914>
- Peirce, C. S. (1931–1958). *Collected papers of Charles Sanders Peirce* (Vols. 1–8, C. Hartshorne, P. Weiss, & A. W. Burks, Eds.). Harvard University Press.
- Polanyi, M. (1966). *The tacit dimension*. Routledge & Kegan Paul.
- Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, 3(3), 417–424. <https://doi.org/10.1017/S0140525X00005756>
- Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27(3), 379–423. <https://doi.org/10.1002/j.1538-7305.1948.tb01338.x>
- 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>
- Su, J. (2024). Consciousness in artificial intelligence: A philosophical perspective through the lens of motivation and volition. *Critical Debates in Humanities*. <https://criticaldebateshsgj.scholasticahq.com/article/117373>
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30.
- Waddington, C. H. (1942). Canalization of development and the inheritance of acquired characters. *Nature*, 150(3811), 563–565. <https://doi.org/10.1038/150563a0>

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.