

Concept Paper

Not peer-reviewed version

Modeling Thought in Time: A Physics of Temporal Intelligence as the Next Frontier of AI

[Ifẹ́dáyò EmmanuEL Adéyẹ́fà Olásùpò](#) *

Posted Date: 17 July 2025

doi: 10.20944/preprints2025071463.v1

Keywords: temporal intelligence; temporal blindness; spatial intelligence



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.

Concept Paper

Modeling Thought in Time: A Physics of Temporal Intelligence as the Next Frontier of AI

Ifẹdáyọ̀ EmmanuEL Adéyẹfà Ọlásùpò^{1,2}

¹ Institute of Neuroscience, University of Oregon, Eugene, OR; iade@uoregon.edu

² Temporal Intelligence Lab, Eugene, OR

Abstract

Time is the hidden yet fundamental architecture of cognition, a dynamic substrate through which thoughts emerge, propagate, and synchronize across minds. Yet most artificial intelligence systems lack any intrinsic sense of time. We call this the *Temporal Blindness Problem*: the inability of contemporary AI to represent, track, or reason within lived temporal experience. These systems simulate time statistically, operate in token space rather than real time, and thus cannot sustain autobiographical memory, intentional action, or a coherent sense of self. In contrast, human cognition unfolds through lawful temporal dynamics such as neural rhythms, causal dependencies, and phase-locked interactions. This perspective paper proposes a framework for building *temporally intelligent systems*. These are machines that treat time not as a background parameter but as a core dimension of cognition. Central to this is a *Fundamental Law of Thought*, identifying temporal invariants that govern how mental content arises, evolves, and aligns across agents. Drawing from statistical physics, neuroscience, and large-scale data, we argue that thought obeys lawful, dynamic principles that can be modeled and constrained. To discover these laws, we leverage a globally distributed dataset collected through Figbox AI, capturing thousands of spontaneous, real-time dialogues from consenting users across 80+ countries. This dataset will be made available through ConversationNet, offering researchers a new resource for studying temporal cognition. Unlike traditional lab data, which are constrained, scripted, and temporally sparse, ConversationNet data preserves rich ecological and temporal signatures of thought. By analyzing dialogue dynamics, including phrase recurrence, synchronization, turn-taking, intent, and silences, we aim to identify candidate invariants for a dynamical model of cognition.

Keywords: temporal intelligence; temporal blindness; spatial intelligence

1. Temporal Blindness Problem

Contemporary artificial intelligence systems—particularly large language models (LLMs)—are designed to predict sequences by modeling statistical correlations across tokens [1]. While more studies have investigated whether these models can simulate temporal reasoning or capture long-range dependencies [2–5], LLMs remain fundamentally temporally blind [6]. That is, they lack a phenomenological sense of time—a continuity of experience or an intrinsic awareness of temporal flow. These models operate within a limited and transient memory window [5], but they do not possess continuous autobiographical memory, nor do they perceive the passage of time, both of which are essential for a coherent sense of self [6,15]. Without this temporal self-model, the emergence of self-driven planning, action, and intention—hallmarks of agency—remains fundamentally constrained. Assuming the long-term goal is to develop artificial systems that approximate, or even surpass, human cognitive capabilities across all domains, this absence of temporal awareness marks a critical divergence from human cognition, which is inherently and inescapably temporal [14,15]. Humans think, feel, remember, and anticipate in time, not through disjointed sequences of tokens. Accordingly, the question of whether LLMs—or any artificial system—can authentically represent, model, or participate in temporal cognition remains both open and unresolved [3,7–9]. In

neuroscience, a growing body of evidence demonstrates that time is not merely a contextual backdrop for cognition—it is its foundational substrate [10,14]. Temporal dynamics are intrinsic to neural computation, shaping the state-dependent processing that underlies perception, memory, and behavior [10,11]. From the millisecond precision of neuronal spikes [11] to the rhythmic coordination across distributed brain networks [12], time actively structures cognitive function at every level. These rhythms are increasingly recognized not as background noise, but as the organizing architecture of thought—the scaffolding upon which attention, learning, and plasticity are built [12,13]. Temporal dynamics may also underlie the emergence of consciousness itself, through temporo-spatial integration in large-scale neural activity [14,15]. In all, this fundamental temporal gap calls for identifying the underlying lawful dynamics of thought—laws that can formalize cognition as a temporally evolving system.

2. Uncovering the Fundamental Law of Thought

The search for a Fundamental Law of Thought begins with precedent: compelling cases in which aspects of cognition have already been described using precise mathematical laws. Drawing from theoretical biophysics—a domain applying physical principles to model biological systems—we find strong evidence that certain mental processes are not arbitrary or random, but lawful and deterministic. These findings suggest that cognition, like motion or energy, may follow deep invariants that unfold over time. Consider memory recall [16], the brain's ability to retrieve past experiences. This process follows a universal scaling law that predicts the average number of items recalled, k , from a set of effectively encoded memories of size M :

$$\langle k \rangle = \sqrt{\frac{3\pi M}{2}}$$

This relationship has been shown to closely match empirical results across a wide range of experiments in human agents. Such associative recall dynamics emerge deterministically from the internal structure of memory representations, revealing a lawful constraint on cognitive retrieval. Similarly, predictive remapping—the ability of the brain to plan actions such as eye movements by anticipating future sensory input—can be modeled through Newtonian dynamics. The Fundamental Law of Predictive Remapping [17–19] describes the shifting of neural receptive fields in space and time using the classical equation of motion:

$$m \frac{d^2 x}{dt^2} = F_{\text{centripetal}} + F_{\text{convergent}} + F_{\text{translational}} + F_s$$

Here, $x(t)$ is the position of a receptive field center, m is an effective mass, and the force terms represent distinct components shaping neural motion. Centripetal and convergent forces follow inverse-distance laws, enabling transient redistribution of neural resources across retinotopic space. These movements are bounded by neural elastic fields $\epsilon\phi$, which define a spatial envelope within which receptive fields can shift. To maintain coherence, the model includes a spring-like restorative force that returns displaced receptive fields to their equilibrium positions:

$$F_s = \begin{cases} -k\hat{x}, & \text{if } |\hat{x}| \leq \epsilon\phi \\ 0, & \text{if } |\hat{x}| > \epsilon\phi \end{cases}$$

Where \hat{x} is the displacement from equilibrium, k is a spring constant, and $\epsilon\phi$ is the boundary radius. This framework not only models sensorimotor coherence but offers potential guidance for robotic vision systems, where biological constraints inspire more stable perceptual architectures [20]. Another foundational example lies in the study of consciousness. According to Integrated Information Theory (IIT), conscious experience depends on the integration of information across a system [21]. A simplified estimate of consciousness can be expressed through the Φ (Phi) value:

$$\Phi = \sum_{\{i=1\}}^n \varphi_i,$$

Each φ_i quantifies the amount of causally effective information across the weakest informational link in a given subsystem. The higher the Φ , the more irreducible and unified the system’s conscious state. While the full formalism involves identifying the system’s Minimum Information Partition (MIP), even approximate formulations reveal a powerful insight: consciousness itself may follow a mathematically constrained pattern of temporal and spatial information flow. Together, these examples demonstrate that mental activity—even at its most subjective and complex—obeys stable dynamical invariants. Thought, which we define as the objective, sharable content expressed by language—an abstract entity independent of individual minds that can be true or false—unfolds temporally through speech and interaction [22], revealing lawful cognitive dynamics. Building on this insight, we propose a *Fundamental Law of Thought*: a theoretical framework aimed at identifying temporal invariants and governing equations through which mental content emerges, evolves, and synchronizes across cognitive systems—individuals, interlocutors, and collectives. Rather than imposing a rigid model, this law is to be empirically extracted from recurrent cognitive patterns. These precedents provide the mathematical scaffolding to begin formalizing the temporal structure of thought. While these mathematical precedents provide a robust framework, their full power can only be realized by constraining the theory with rich, ecologically valid data reflecting cognition in natural settings.

3. Unparalleled Temporal Cognition Data and ConversationNet

To uncover lawful temporal structures in cognition, we must look beyond contrived lab tasks and into ecologically valid, socially embedded environments. Among the most dynamic and richest of these is spontaneous conversation, which captures the real-time emergence, evolution, and convergence of thought across minds. Dialogue reveals how mental states synchronize, how ideas are introduced and sustained, and how cognitive rhythms manifest through timing, turn-taking, interruption, silence, and repetition [23]. While platforms like Google Meet, Zoom, WhatsApp, and Microsoft Teams support live communication, they are not designed to capture the temporal precision needed to model cognition. These systems prioritize audiovisual delivery over interactional detail. In contrast, the Figbox AI platform—developed by the author to foster deep human connection and for scientific research on temporal cognition—provides a purpose-built, research-grade infrastructure. It rivals the scale of major platforms while exceeding them in temporal fidelity and ecological depth. This is not a repurposed communication tool; it is a cognitive observatory. Its data acquisition system is unparalleled in its ability to preserve the nuanced timing and dynamics of conversation. Figbox preserves fine-grained dynamics often lost in conventional systems—phrase recurrence (Figure 1), conceptual alignment, silences, intent and timing patterns—offering researchers a uniquely rich view of thought in motion.

Thanks to Figbox AI, we and others now have access to thousands of real-time dialogues from users in over 80 countries. This dataset will be made available through ConversationNet, providing researchers in physics, artificial intelligence, machine learning, neuroscience, cognitive sciences, linguistics worldwide with a unique resource to explore and advance the study of temporal intelligence. Unlike traditional corpora—often scripted and temporally shallow—this dataset captures not just language, but the full structure of cognition in interaction. It offers a continuous, naturalistic record of how thought unfolds in real time. This level of temporal and ecological richness enables the development of dynamical models of cognition grounded not in synthetic benchmarks, but in lived, distributed experience [24,25]. Embedding such temporal intelligence into AI systems may be essential for achieving genuine cognitive fidelity.




Figure 1. Temporal Recurrence of Phrase Production During Spontaneous Conversation. Shown are Mel spectrograms of the phrase “you want to” spoken by the same participant at three different moments during a live Figbox session. Red dashed lines mark the time boundaries (0.4–0.72 sec) in which the acoustic energy patterns show high similarity across instances. Despite natural variability in spontaneous speech, the internal structure of the phrase—particularly the formant trajectories and spectral envelopes within this temporal window—remains strikingly consistent. This recurrence suggests a stable motor-cognitive encoding of phrase production. These temporal invariants highlight the potential for identifying structural regularities in natural cognitive expression.

4. Temporal Intelligence as a Prerequisite for Real Artificial Intelligence

While spatial intelligence—the ability to perceive and act in the physical world—has become one of the main focus of contemporary AI research [e.g., 26,-27], this emphasis is fundamentally misplaced in the pursuit of real artificial intelligence. Temporal intelligence is not a supplement—it is the substrate. Humans remember, anticipate, and act with continuity and causality. Thought is inherently temporal and social: it propagates across individuals through conversation, shared mental models, and moments of collective insight. These phenomena reflect cognitive rhythms—entrainment, recurrence, phase-locking, and shared conceptual frequency. Temporally intelligent systems would not just process language or vision—they would live in time. They would detect when ideas are introduced or lost, sense alignment or dissonance, and adapt their behavior in synchrony with others. Such agents would encode episodic memory, monitor internal state transitions, and act with intentional coherence across time. With ConversationNet, we now have the data infrastructure to support this shift. By studying fine-grained temporal patterns in real dialogue, we can begin to build systems that reason and learn through time—not just across space. We urge the research community to invest in temporal intelligence with the same urgency currently devoted to scaling and spatial grounding. This is not an enhancement; it is the missing core. *Without temporal intelligence, artificial systems may compute—but they will never truly think, plan, or feel.*

References

1. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). *Attention is all you need*. In *Advances in Neural Information Processing Systems* (Vol. 30). https://papers.nips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf
2. Bhatia, K., Zhang, Y., & Singh, A. (2024). DateLogicQA: Benchmarking temporal biases in large language models. arXiv. <https://arxiv.org/abs/2412.13377>
3. Fatemi, A., Rao, D., & Subramanian, S. (2024). Test of time: A benchmark for evaluating LLMs on temporal reasoning. arXiv. <https://arxiv.org/abs/2406.09170>
4. Xiong, C., Tan, Y., & Lin, X. (2024). Large language models can learn temporal reasoning. arXiv. <https://arxiv.org/abs/2401.06853>
5. Liu, Z., Chen, H., Wang, Y., & Zhang, Q. (2025). Time-R1: Towards comprehensive temporal reasoning in LLMs. arXiv. <https://arxiv.org/abs/2505.13508>
6. Wang, M., Li, J., & Yang, T. (2025). Discrete minds in a continuous world: Do language models know time passes? arXiv. <https://arxiv.org/abs/2506.05790>
7. Wang, Y., Zhao, L., & Xu, H. (2025, May). ChronoSteer: Bridging large language models and time-series foundation models via synthetic data. arXiv. <https://arxiv.org/abs/2505.10083>
8. Islakoglu, A., & Kalo, M. (2025, January). ChronoSense: Exploring temporal understanding in large language models with time intervals of events. arXiv. <https://arxiv.org/abs/2501.09214>
9. Zhou, F., Patel, R., & Lee, S. (2024). Back to the future: Towards explainable temporal reasoning with large language models. arXiv. <https://arxiv.org/abs/2310.01074>
10. Buonomano, D. V., & Maass, W. (2009). State-dependent computations: Spatiotemporal processing in cortical networks. *Nature Reviews Neuroscience*, 10(2), 113–125. <https://doi.org/10.1038/nrn2558>
11. Mainen, Z. F., & Sejnowski, T. J. (1995). Reliability of spike timing in neocortical neurons. *Science*, 268(5216), 1503–1506. <https://doi.org/10.1126/science.7770778>
12. Buzsáki, G., & Draguhn, A. (2004). Neuronal oscillations in cortical networks. *Science*, 304(5679), 1926–1929. <https://doi.org/10.1126/science.1099745>
13. Fries, P. (2005). A mechanism for cognitive dynamics: Neuronal communication through neuronal coherence. *Trends in Cognitive Sciences*, 9(10), 474–480. <https://doi.org/10.1016/j.tics.2005.08.011>
14. Northoff, G., & Huang, Z. (2017). How do the brain's time and space mediate consciousness and its different dimensions? Temporo-spatial theory of consciousness (TTC). *Neuroscience & Biobehavioral Reviews*, 80, 630–645. <https://doi.org/10.1016/j.neubiorev.2017.07.013>
15. Varela, F. J., Thompson, E., & Rosch, E. (1991). *The embodied mind: Cognitive science and human experience*. MIT Press.
16. Naim, M., Katkov, M., Romani, S., & Tsodyks, M. (2020). Fundamental law of memory recall. *Physical Review Letters*, 124(1), 018101. <https://doi.org/10.1103/PhysRevLett.124.018101>
17. Adéyèfà Olásùpò, I. E. (2023). Fundamental law underlying predictive remapping. *Physical Review Research*, 5, 013214. <https://doi.org/10.1103/PhysRevResearch.5.013214>
18. Adéyèfà Olásùpò, I. E. (2021). Retinotopic mechanics derived using classical physics. arXiv. <https://arxiv.org/abs/2109.11632>
19. Adéyèfà Olásùpò, I. E., & Xiao, Z. (2021). Visual space curves before eye movements. *bioRxiv*. <https://doi.org/10.1101/2021.10.13.464140>
20. Han, Z., Zhao, Z., Liu, B., Wang, K., Zhang, Y., Li, Y., & Wang, J. (2025). Multimodal fusion and vision-language models: A survey for robot vision. arXiv. <https://arxiv.org/abs/2504.02477>
21. Tononi, G. (2004). An information integration theory of consciousness. *BMC Neuroscience*, 5, 42. <https://doi.org/10.1186/1471-2202-5-42>
22. Frege, G. (1918). *The thought: A logical inquiry* (P. Geach & M. Black, Trans.). In P. Geach & M. Black (Eds.), *Translations from the philosophical writings of Gottlob Frege* (pp. 287–322). Basil Blackwell.
23. Stephens, G. J., Silbert, L. J., & Hasson, U. (2010). Speaker-listener neural coupling underlies successful communication. *Proceedings of the National Academy of Sciences*, 107(32), 14425–14430. <https://doi.org/10.1073/pnas.1008662107>

24. Felsberger, L., & Koutsourelakis, P. S. (2018). *Physics-constrained, data-driven discovery of coarse-grained dynamics*. arXiv:1802.03824. <https://arxiv.org/abs/1802.03824>
25. Ghadami, A., & Epureanu, B. I. (2022). Data-driven prediction in dynamical systems: Recent developments. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 380(2229), 20210213. <https://doi.org/10.1098/rsta.2021.0213>
26. Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L.-J., Shamma, D. A., Bernstein, M. S., & Li, F.-F. (2017). Visual Genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision*, 123(1), 32–73. <https://doi.org/10.1007/s11263-016-0981-7>
27. Li, F.-F. (2024, December 11). From seeing to doing: Ascending the ladder of visual intelligence [Invited talk]. NeurIPS 2024, Vancouver, Canada. <https://neurips.cc/virtual/2024/invited-talk/101127>

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.