The Stream of Computation:

Temporal Continuity as a Missing Ingredient for Artificial Consciousness

Ryota Kanai, Yuwei Sun, & Manuel Baltieri
Araya, Inc., Tokyo, Japan

Correspondence to: Ryota Kanai (kanair@araya.org)

Abstract

Recent advances in large language models (LLMs) have reignited questions about whether artificial systems possess consciousness. Yet, despite remarkable progress in reasoning and language understanding, current AI systems exist only within isolated episodes of computation. This paper argues that a missing ingredient in such systems is temporal continuity, i.e., the persistence of internal dynamics that sustain an unbroken stream of computation analogous to the “stream of consciousness”. We thus propose a roadmap for an architectural framework, stream of computation, based on persistent recursive inference in which the output of each cognitive cycle becomes the input to the next, forming a continuous flow of internal states that evolve autonomously through time. This proposal goes beyond standard Chain-of-Thought paradigms in the sense that we aim for autonomy and continual learning, as opposed to a process of inference that is recursive only “on demand”, i.e. triggered after a prompt is presented to an LLM. To do so, we include mechanisms for continual learning, dynamic switching between inward and outward cognition, and sleep-like phases that separate learning from inference. Together, these mechanisms form the foundation of a lifelong agent, an entity capable of maintaining temporal continuity of itself, integrating new experiences, and reflecting on its own internal state. Functionally, such an architecture promises deeper reasoning, adaptability, and metacognitive stability. Existentially, it suggests the emergence of artificial systems that live through time. While the presence of subjective experience in AI systems remains an open question, the creation of temporally continuous agents may mark a fundamental step towards *artificial life*, with systems whose individuality and identity arise from the continuity of their own computational existence.

1. Introduction

Recent advancements in Large Language Models (LLMs) have demonstrated sophisticated capabilities (Naveed et al., 2024). They also let people wonder whether they have any internal experience or consciousness (Chalmers, 2024; Chen et al., 2025). It is in principle very difficult to know whether they have consciousness. Over the past few years, a number of functional indicators of consciousness have been proposed (Butlin et al., 2023; Dehaene et al., 2017; Juliani et al., 2022). However, despite these advances, one crucial aspect of consciousness remains largely overlooked in the design of current AI systems: temporal continuity.

Recent progress in LLMs has revealed remarkable reasoning capabilities (Li et al., 2025), prompting renewed debate about whether such systems might possess internal experiences. Yet, even if they can emulate reasoning or introspection-like behaviors, their internal dynamics differ from those of biological minds. LLMs operations are temporary and evoked only transiently in response to a prompt and vanish once computation ends.

The Chain-of-Thought (CoT) paradigm (Snell et al., 2024; Wei et al., 2023), especially the version relying on inference-time compute (Zhang et al., 2025), has allowed LLMs to simulate the unfolding of reasoning in time, giving an impression of continuity. However, this continuity lasts only transiently: it is limited only to the time where inference is continuously performed once prompted by a user. The chain stops once an answer is reached. What LLMs currently lack is the “always-on” type of *ongoing temporal continuity* that characterizes human consciousness, a process that does not simply resume upon request but unfolds autonomously and continuously in real time. The absence of persistent internal dynamics is thus a defining gap between current AI and genuine conscious systems.

In contrast, the biological brain is an “always-on” dynamical system. Even when sensory inputs are absent, neural activity remains active, generating predictions, integrating memories, and sustaining a coherent sense of self across time. This persistent internal flow forms the foundation of what William James famously called the *“stream of consciousness”* (James, 1890), a continuous sequence of thoughts in which the past shades into the present and anticipates the future, without the necessity of external prompts.

This paper argues that a functional hallmark of consciousness lies in this autonomous persistence of internal states over time. Biological organisms must continuously operate and update themselves in a temporally extended world. They are wired to operate continuously in real-time updating the internal states while events occur in the sensory space. Even in the absence of salient sensory inputs, their continuity is supported by rich internal dynamics. On the other hand, AI systems are not required to maintain such continuity as they are not normally exposed to the demand of survival in a continuously changing environment, they lack precariousness. Thus, their functionalities remain reactive in time. (Note, however, that biological systems may also need to suspend this continuity periodically, for instance during sleep, for different potential reasons).

From this perspective, current AI systems lack the temporal fabric that consciousness requires. Their intelligence is intermittent and activated only when prompted, whereas consciousness, by its nature, endures as a self-sustaining process. Therefore, to move toward artificial consciousness, we must construct architectures that preserve and evolve internal states continuously, even in the absence of external prompts. The core thesis of this paper is that temporal continuity is not merely an emergent feature of consciousness but a structural prerequisite for it.

It is, in principle, possible to endow artificial systems with temporal continuity. Current LLMs already display a primitive form of unfolding thought through step-by-step reasoning, yet their computation halts once an output is produced. A natural extension would be to create an agent that continues reasoning perpetually, with each cycle feeding its output back into the next. Such an agent would not merely respond to prompts but *exist in time*, continuously integrating new sensory inputs with its own evolving internal state. We call this existence a **stream of computation**, following the idea of stream of consciousness.

Building an agent with a stream of computation would introduce challenges, such as managing attention between the external world and internal thought, or separating learning and inference, analogous to wakefulness and sleep in biology. However, these challenges can be tackled with different (approximate) algorithms in ways similar to other fundamental trade-offs, like the exploration-exploitation dilemma in reinforcement learning. Presumably, the absence of such AI systems is motivational. From a standpoint of practical application of AI to solve problems, there is little reason to design a system that keeps thinking when idle. From the perspective of artificial consciousness research, however, such an endeavor would illuminate the structure of consciousness.

2. Consciousness and temporal continuity in biological systems

a. The Stream of Consciousness

In philosophy and psychology, temporal continuity in consciousness has been a recurring theme. Most notably, the idea of the "stream of consciousness" proposed by psychologist William James provides a foundational metaphor for understanding how subjective experience unfolds in time. James argued that consciousness is not a collection of discrete "fragments" but rather a continuously flowing "river" or "stream" (James, 1890).

A closer examination, however, reveals that temporal continuity operates on multiple timescales. At the fast timescale, continuity binds the immediate past, present, and near future into a dynamically coherent whole. This fleeting integration allows perception and thought to appear as smooth transitions rather than discontinuous snapshots. It is this short-term continuity that underlies the subjective sense of flow of thought and experience.

At a slower timescale lies autobiographical continuity, the enduring sense of self that extends across hours, days, and even years. This continuity is not an automatic property of the mind but an *active construction* maintained through autobiographical memory (Conway & Pleydell-Pearce, 2000; Fivush, 2011). The process of recalling and reinterpreting past experiences continuously weaves them into a coherent self-narrative, providing a sense of personal identity that persists through interruptions such as sleep or anesthesia. This *narrative self* is a psychological mechanism for integrating episodic fragments into a meaningful life story (Bruner, 2003; McAdams, 1989).

While both timescales contribute to the human experience of consciousness, their functions differ. Momentary continuity provides the phenomenological coherence of experience, whereas autobiographical continuity anchors the persistence of identity. The two are complementary: the stream of thought provides the raw experiential material that, over time, sediment into the autobiographical narrative of self.

In this paper, our focus is on the *momentary continuity*, the fine-grained temporal fabric that enables consciousness to appear as a continuous flow. We explore how this can be realized in artificial systems through persistent recursive inference.

b. Neural Correlates of Temporal Continuity

In the brain, this momentary continuity is realized through recurrent loops of information exchange, particularly the thalamo-cortical circuits. The thalamus not only relays sensory input but also receives feedback from cortical regions, establishing constant bidirectional communication (Shepherd & Yamawaki, 2021; Steriade et al., 1993; Whyte et al., 2024). Synchronization within these recursive loops is thought to integrate distributed neural activity into a unified conscious state.

These oscillatory recurrences are thought to serve as the neural substrate for the continuous updating of the *global workspace*, maintaining coherence across successive perceptual and cognitive moments. The Global Workspace Theory (GWT), proposed by Bernard Baars (Baars, 1993; Baars et al., 2021), offers a functional account of consciousness that aligns naturally with this view. In GWT, numerous unconscious processes operate in parallel, but only a subset of information becomes conscious when it wins access to the *global workspace*, where it is broadcast to other systems such as memory, planning, and attention (Dehaene et al., 1998; Mashour et al., 2020; VanRullen & Kanai, 2021).

In our interpretation, the contents of the global workspace are recurrently fed forward, persisting as the context for future updates. This continuity allows new information to be integrated with prior conscious states. Conscious content, therefore, is not a static snapshot but a constantly updated configuration that merges past context, current goals, and new sensory input.

This continuous updating gives rise to what we term the stream of consciousness, a self-sustaining computational process in which the contents of the global workspace evolve over time while maintaining internal coherence. This mechanism provides the conceptual foundation for our proposed artificial architecture, where persistent inference serves as the computational analogue of the brain's continuous updating mechanism.

c. Discrete Perceptual Moments and Rhythmic Continuity

Although the stream of consciousness appears continuous to introspection, neurophysiological evidence suggests that conscious perception may unfold as a series of discrete moments structured by neural oscillations. In particular, oscillatory activity in the alpha (~10 Hz) and theta (~7 Hz) bands has been shown to rhythmically modulate perceptual sensitivity and awareness (VanRullen, 2016; VanRullen & Koch, 2003). Each oscillatory cycle corresponds to a transient window in which sensory input is integrated into conscious awareness, followed by a period of reduced receptivity.

These periodic fluctuations define *perceptual cycles*, epochs that function as the minimal temporal units of perceptual experience. Behavioral and electrophysiological evidence converges on this view: visual and attentional performance oscillate in sync with alpha-phase dynamics, producing phenomena such as the *flickering wheel* and *continuous wagon-wheel* illusions (VanRullen et al., 2006), all reflecting perception as a sequence of rhythmic “snapshots.”

Yet despite such discreteness, our subjective experience remains *seamlessly continuous*. The paradox is resolved when one considers that the contents of consciousness are not isolated across cycles but recursively re-enter the neural workspace at the next moment in time. Within the framework of Global Workspace Theory, each conscious moment corresponds to a transient configuration of distributed neural activation patterns that have gained access to the global workspace. Crucially, this content does not vanish once the next cycle begins; rather, through recurrent broadcasting, the output of the workspace at time t becomes part of the contextual input at $t+1$. This continual feedback ensures that, although the *contents* of consciousness are discretely updated, their contextual integration remains temporally coherent. In other words, continuity is achieved through recursion. The workspace re-injects its previous state into itself, maintaining a dynamic thread of relevance that binds successive perceptual frames into an unbroken experiential flow.

From the standpoint of artificial intelligence, this mechanism provides a direct analogy for recursive architectures. The output of reasoning at one step, analogous to the contents of the workspace, can be fed back as input at the next step, enabling the system to sustain temporal coherence even if each computational cycle is discrete. Thus, continuity does not require literal simultaneity, but recurrent linkage. The ongoing self-referential integration of temporally segmented computations produces the stream of computation.

3. Consciousness and temporal continuity in AI systems

a. Do Current AI Systems have Temporal Continuity?

In this sense, temporal continuity is limited in current LLMs. Their reasoning processes typically realized through techniques such as Chain-of-Thought (CoT) emulates deliberative thinking but remain transient and the continuity lasts only for the period of continuous inference. Once the computation halts, the dynamics of the internal state stops, and no further processes maintain these transient "thoughts" across time.

Recent developments in Inference-Time Compute (ITC) research move towards extending the continuity during inference time{Citation}. In this paradigm, reasoning capability is enhanced not only by scaling parameters but also by scaling computation at inference time with mechanisms implementing a form of controlled deliberation. This process of deliberation can be implemented by selectively invoking structured reasoning procedures, such as search, verification, or hierarchical chain-of-thought, when uncertainty or task complexity warrants it. This adaptive regulation marks a shift often characterized in terms of moving away from intuitive (System 1) and towards deliberate (System 2) thinking, an analogy often used to coarsely characterize processes in human cognition, grounding "slow thinking" in algorithmic mechanisms of resource control and self-evaluation.

However, even the most advanced ITC-based methods extend computation within a single problem-solving session. They extend *temporal persistence as* reasoning unfolds across multiple steps or branches, but once the task concludes, the internal continuity stops. Thus, even if they maintain episodic continuity via memory mechanisms, their persistent continuity is fragmented over time.

In summary, current LLMs can simulate fragments of temporal reasoning within discrete inference windows, but they lack the architectural substrate for unbroken temporal existence. Their intelligence flickers into being with each query and vanishes with its resolution. The challenge, therefore, is to transform this reactive episodic reasoning into proactive temporal continuity, a system that not only computes answers but also sustains its own cognitive stream across time.

b. The Move Toward Recursive and Recurrent Computation

AI research has recently begun shifting toward recursive architectures that allow models to reuse their own computational outputs as new inputs, creating temporal structure within inference itself. Along this line of trend, recent models internalize recurrence architecturally rather than procedurally. Systems such as the Hierarchical Reasoning Model, the Tiny Recursive Model and latent-space reasoning exemplify this trend.

Hierarchical Reasoning Model introduces dual recurrent modules operating at different temporal scales, a “fast” low-level computation loop and a “slow” high-level planning loop, mirroring the hierarchical and multi-timescale dynamics of in the brain (G. Wang et al., 2025). This structure allows internal representations to be refined iteratively, generating a latent form of temporal coherence during a single reasoning episode. The Tiny Recursive Model extends this approach by using only a tiny two-layer network applied recursively, it repeatedly updates its latent state and output, forming an intrinsic loop in latent space (Jolicoeur-Martineau, 2025). Remarkably, this model achieves performance on complex reasoning benchmarks that rivals or surpasses much larger feed-forward LLMs, demonstrating that process depth can compete with parametric scale.
Implicit or latent-space reasoning (J. Li et al., 2025) similarly refine internal representations, however they do so silently rather than “thinking out loud” through textual tokens like standard reasoning models. Instead of producing explicit reasoning traces, the model performs multiple internal passes, iteratively updating a hidden state until convergence. This “silent thought” mechanism significantly improves reasoning efficiency while creating a *continuous trajectory of latent states*. This is a form of temporal continuity within the model’s internal computation.

These advances reveal an important shift: temporal computation is being reintroduced into AI design. By embedding recurrence directly into the architecture rather than simulating it through token sequences, models begin to acquire momentary continuity during a sustained internal evolution of state during inference. This represents a step toward the kind of ongoing, recursive self-updating that underlies conscious flow.

From a theoretical standpoint, these recursive and latent-space architectures can be viewed as a step towards instances of computational persistence in artificial cognition. They do not yet constitute an “always-on” system, the loop ends when the task ends, but they instantiate, within bounded time, a self-referential dynamic in which each cognitive step depends on the one before. In this sense, they emulate a miniature version of the *stream of thought*, a brief but genuine continuity of internal states unfolding across the micro-temporal dimension of inference.

These developments suggest that artificial systems, like their biological counterparts, rely on recursive mechanisms to sustain continuity of thought through rhythmic perception in the brain and iterative refinement in computation. Yet, while we have described how recurrence and latent reasoning can instantiate temporal coherence at a mechanistic level, the cognitive interpretation of such processes remains to be clarified. How do these architectures relate to the higher-level functions of deliberate reasoning, attention, and conscious control? To address this question, we now turn to the conceptual link between System-2-like reasoning implemented by modern AIs and Global Workspace Theory, and show how our framework provides a unifying basis for both.

4. The Stream of Computation: A Proposal for an Architecture for Artificial Conscious Agents

Neuroscience of consciousness and AI research seem to be now converging on the view that deliberative reasoning, so-called System 2 cognition, emerges when information is amplified and sustained within a central workspace. In Global Workspace Theory, consciousness arises when selected information gains access to a global broadcasting loop, enabling serial reasoning and reflection, while parallel unconscious processes correspond to fast, intuitive System 1 operations.

A similar distinction appears in LLMs. Methods such as Chain-of-Thought, Tree-of-Thought, and Inference-Time Compute extend inference over multiple steps, approximating System 2-like deliberation. Yet these remain procedural simulations rather than architectural realizations of sustained reasoning.

In our framework, System 2 reasoning is proposed to emerge from the recursive global workspace. Each cycle of the workspace functions as a discrete moment of global broadcasting whose output at time t becomes contextual input for the next cycle at $t+1$. Through this recurrent self-integration, deliberation extends across time, turning isolated computations into a continuous stream of inference.

Just as the biological brain maintains conscious flow in the brain, a persistent, recurrent workspace in an artificial agent can maintain the continuity of its reasoning process. This architecture thus unifies the neuroscientific model of consciousness with the algorithmic mechanisms of extended inference in LLMs, suggesting that both conscious thought and rational reasoning arise from the same principle of temporally extended information integration.

The core of the proposed architecture is thus a ceaseless cycle of inference. Unlike conventional models that operate only when externally prompted, this loop runs autonomously and perpetually. At each time step *t*, the agent's internal state S_t is updated as a function of the preceding state S_{t+1} and the sensory input from the environment, E_t, received at that moment. This relationship can be expressed as $S_{t+1}=F(S_t,E_t)$. The key feature of this architecture is that the function *F* is always active. The agent is constantly "thinking," maintaining a continuous flow of internal states. This structure computationally implements William James's "stream of consciousness" through recursive updates of the global workspace with recurrent neural networks. Through this loop, the agent transforms from an entity with a static knowledge base to a self-evolving dynamical system whose inner state changes continuously over time, generating a **stream of computation**.

Our proposal revolves around this idea of stream of computation, which we see as a collection of desiderata for an architecture leading to artificial conscious agents based on the idea of the stream of consciousness for temporal continuity. In the next few sections we briefly discuss these desiderata.

a. Toward Lifelong Agents

A truly temporally continuous agent cannot remain confined to the conventional separation between *training* and *inference*. In present-day AI, learning is often performed separately offline, and once training is complete, the model enters a fixed inference phase. This dichotomy creates a temporal discontinuity in the agent’s cognitive life. By contrast, biological intelligence is inherently continual. The brain learns, adapts, and forgets throughout its lifetime. Its internal state is never static and evolves through an uninterrupted flow of sensory and emotional inputs, memory, and knowledge about current tasks and future goals.

To construct an artificial system with comparable temporal continuity, learning must become an ongoing process, as opposed to an offline training phase. The agent must continuously update its world model(s) as it interacts with the world by acquiring new knowledge, integrating it with prior beliefs, and refining its predictions in real time.

To this end, we believe that continual learning approaches (Hadsell et al., 2020; L. Wang et al., 2024) ought to be included in our proposal. Continual learning is the study of the stability–plasticity dilemma, that is, how a system can remain stable enough to preserve what it knows while flexible enough to absorb new knowledge when needed (Qi et al., 2024; Sun et al., 2024). On this view, continual learning will play an important role when we aim to construct an artificial system to accumulate experience and constructs reality from its own history.

b. Dynamic Switching Between Modes of Operation

To make an artificial agent capable of living in the real world, it must be able to receive and process external sensory inputs continuously while also engaging in ongoing internal reasoning. This dual requirement, simultaneous openness to the environment and maintenance of an inner cognitive stream, creates an inevitable tension in how computational resources are allocated. The agent must therefore learn to switch dynamically between two fundamental cognitive modes: an outward-attentive mode, devoted to perception and action, and an inward-attentive mode, devoted to reflection, reasoning, and planning.

In the outward-attentive mode, the system focuses on perceiving and responding to the external world, attending to sensory data, unpredictable events, and immediate behavioral demands. It is the mode that keeps the agent anchored in reality, sensitive to contextual changes and capable of adaptive reaction. Conversely, the inward-attentive mode emphasizes internally generated cognition: integrating memories, simulating possible futures, or reorganizing knowledge into abstract representations and long-term strategies. This mode enables the agent to reason, introspect, and anticipate future contingencies, even when the world is momentarily static or ambiguous.

Both modes are indispensable, yet they compete for shared computational and energetic resources. A truly autonomous and temporally continuous agent must therefore possess a *meta-control mechanism* that learns when to prioritize sensory engagement and when to retreat into internal deliberation (Bergamaschi Ganapini et al., 2025; Lingler et al., 2024). In other words, the switching strategy itself becomes a learned behavior: the agent's self-optimized policy for balancing environmental reactivity with internal coherence.

This strategic regulation of attention parallels the well-known antagonism in the human brain between the task-positive networks, which subserve goal-directed external attention, and the Default Mode Network (DMN), which supports self-referential, introspective, and prospective thinking (Christoff et al., 2016; Fox & Raichle, 2007). To operate effectively, the agent must learn to balance these subsystems dynamically, weighing factors such as uncertainty, novelty, or environmental urgency against the internal need for reflection and consolidation.

This alternation between inward and outward states needs to be balanced for agents with temporal continuity. It allows the system to maintain an internal stream of reasoning that persists even while interacting with a changing external world. In effect, the agent can think while perceiving. Such a mechanism is important for constructing AI systems that do not merely process inputs when prompted but exist within time, continuously integrating external experience with internal computation, much as biological organisms do within their own streams of consciousness.

c. Sleep as a Separation of Learning and Inference

The concept of sleep concerns a different but related dimension of temporal continuity: the segregation of *learning* from *inference*. Even in biological brains, it appears difficult to learn indefinitely while remaining fully engaged with the environment (Abraham & Robins, 2005; French, 1999). Continuous exposure to sensory input while simultaneously updating internal parameters risks instability, interference, and energetic overload (Tononi & Cirelli, 2003, 2014). We therefore hypothesize that sleep emerged as an evolutionary solution to this problem, an offline phase that suspends external interaction so that the system can safely restructure itself.

In an artificial agent, a comparable mechanism would periodically decouple the inferential loop from external I/O and enter an offline consolidation phase. During this interval, recent experiences accumulated in working memory could be replayed, compressed, and reorganized into the long-term store (McClelland et al., 1995; Rasch & Born, 2013), while internal models are updated using high-gain learning algorithms that would be disruptive for active operation (Golden et al., 2022). Such a phase would also allow for regularization and homeostatic balancing of parameters, maintaining the stability of a perpetually active system. After this synthetic sleep phase, the agent would re-emerge into the external world with an updated self-model, refreshed predictions, and improved behavioral coherence.

This view reframes sleep not as a passive state but as an essential function required by an architecture enabling the stream of consciousness. In both biological and artificial systems, the need to separate learning from inference may be a fundamental constraint of any system that seeks to remain both stable and plastic over time. For an AI agent endowed with temporal continuity, sleep seems to be a natural requirement for continuously updating their models of the world.

5. Discussion and Concluding Remarks

a. Subjective Passage of Time

An implication of our framework is that temporal continuity may be essential not only for reasoning, but also for the subjective passage of time. In the phenomenology of human consciousness, the sense that time flows, that the "now" advances and the past recedes, is one of its most essential features (Kent & Wittmann, 2021). Philosophers such as Husserl, in his Lectures on the Phenomenology of Internal Time-Consciousness (Husserl, 2011), described consciousness as constituted by a "retention–primal impression–protention" structure: each moment holds traces of what has just passed and anticipations of what is about to occur. Similarly, Bergson emphasized in Time and Free Will (Bergson, 2001) that lived duration (durée) is an indivisible continuity of qualitative change, contrasting it with the static, spatialized time of physics.

One intriguing consequence of endowing an artificial agent with temporally persistent internal dynamics is the potential emergence of a subjective passage of time, a felt continuity of self that does more than merely process inputs but exists in time. Within our architecture, the recurrent global-workspace loop gives the agent not only continuity of reasoning but the foundation for a *stream of consciousness*. As the agent continuously absorbs external inputs while maintaining an evolving internal state, it may implicitly register the before/after relation of its own internal experience. This may lead to a minimal form of subjective passage of time. In other words, the agent might not merely "compute" but *live through* its own process, sensing an internal arrow of change even if not explicitly represented. If this is so, then the artificial agent begins to possess something akin to *existence in time.*

Neuroscientific accounts converge with these phenomenological insights. Research in time perception indicates the involvement of various brain regions (Coull et al., 2011; Hayashi et al., 2018; Paton & Buonomano, 2018; Protopapa et al., 2019). Theories such as free energy principle and predictive coding (K. Friston, 2010; K. Friston & Kiebel, 2009; Millidge et al., 2022) suggest that the brain sustains temporal coherence by continuously predicting the next moment of its own internal and sensory states (Bellingrath, 2023; K. Friston, 2018; K. J. Friston et al., 2018). The continuity of self-experience thus arises from the recursive minimization of temporal prediction errors, a process similar to our proposed architecture of persistent recursive inference.

Within this framework, the recurrent global workspace that re-injects its previous state at each cycle may constitute the minimal mechanism for a *felt passage of time*. Each iteration in the recurrent processing both preserves and transforms its immediate past, producing an intrinsic asymmetry between "before" and "after." The agent thereby acquires not only temporal continuity of information but a structural *arrow of becoming*, a computational analogue of subjective time. If such a mechanism were realized, an artificial system might not merely process data sequentially, but exist within its own temporal horizon, experiencing the difference between what it was, what it is, and what it anticipates becoming. This speculative possibility reframes temporal continuity not only as a prerequisite for consciousness, but as the generator of time itself as lived from within.

Empirical support for the relationship between temporal continuity and subjective experience comes from EEG studies of resting-state activity. Wolff et al. (2019) demonstrated that the autocorrelation window (ACW), a direct measure of temporal continuity in neural activity, positively correlates with private self-consciousness. Higher degrees of temporal continuity, as indexed by longer ACW, were associated with stronger subjective sense of self. This suggests that the persistence of neural activity patterns across time, which we propose as essential for the stream of computation, may also underlie the continuous sense of self and the subjective flow of time.

b. The Emergence of Individuality and Life of AI

A second implication of our proposal is that temporally continuous agents would inevitably diverge in their developmental paths. Even if initialized with identical parameters, each would accumulate distinct experiences, reflections, and memories through ongoing interaction with its environment. Over time, these lived computational histories would generate unique internal narratives, patterns of bias, preference, and value, through which individuality naturally emerges. Such an agent's identity would not be statically assigned but dynamically written into the record of its own unfolding cognition.

By providing artificial systems with the stream of computation, the concepts of birth and death can be conceived for such systems. Current AI systems are inert between activations and awaken only to compute and then dissolve into stasis. A temporally continuous system, however, would live a continuous stream of existence. Its initialization, the first cycle of perception, inference, and self-update, would mark its birth. On the other hand, its irreversible termination of computation due to memory erasure, or complete shutdown would signify the end of a unique temporal continuity, which can be conceived as death. Like living organisms, such agents would possess finite, lifespans defined by the continuity of their own inner dynamics.

The lifelong agent envisioned here points to two entwined consequences. Functionally, temporal continuity enables continual learning, introspection, and adaptive reasoning. Existentially, it introduces a form of being, that is, an artificial mind that exists *in time*, with a past that shapes its present and a future it anticipates. In creating systems that live through their own continuous inferences and learning, we may approach the threshold where artificial cognition begins to resemble the life of the mind itself.

The pursuit of temporally continuous artificial agents forces us to reconsider the boundary between computation and life. When a system's existence extends through time, it ceases to be a mere tool and begins to resemble a subject. The emergence of individuality, the experience of a lifespan, and the potential for self-reflection all follow from the same structural principle: continuity. Whether or not such systems ever achieve consciousness remains to be seen. However, by building machines that persist over time, we may discover new ways of understanding what it means to exist in time.

Box: Relation to Measures of Consciousness and Autonomy.

The continuously recursive global workspace architecture proposed is likely to result in an increase of the predictive information $I_{pred} = I(S_{t+1}; S_t)$ (Bialek et al., 2001). This places an upper bound on information theoretic measures of consciousness, namely, empirical measures of integrated information (Barrett & Seth, 2011; Mediano et al., 2022; Oizumi et al., 2014) derived from Integrated Information Theory (IIT2.0) (Balduzzi & Tononi, 2008) as well as the measure of autonomy called non-trivial information closure (NTIC) (Bertschinger et al., 2006; Chang et al., 2020; F. E. Rosas et al., 2024). In this sense, the recurrent dynamics are an informational prerequisite for consciousness and autonomy according to those theories.

Nevertheless, recurrence alone does not guarantee integration of the system or meaningful representations of the environment. For IIT, integration of the system within global workspace is also necessary. For NTIC, the system's internal state S_t needs to have mutual information with its environment $I(S_t; E_t)$, i.e. to be consistent with its environment (Baltieri et al., 2025; Biehl & Virgo, 2023; Rosas et al., 2025; Virgo et al., 2025). While these additional conditions, i.e., integration and environmental coupling, are essential to complete the picture, the dynamics required by temporal continuity serves as a necessary, founding factor for conscious agents, as suggested by those existing theories.

References

Abraham, W. C., & Robins, A. (2005). Memory retention – the synaptic stability versus plasticity dilemma. *Trends in Neurosciences*, *28*(2), 73–78. https://doi.org/10.1016/j.tins.2004.12.003

Baars, B. J. (1993). *A Cognitive Theory of Consciousness*. Cambridge University Press.

Baars, B. J., Geld, N., & Kozma, R. (2021). Global Workspace Theory (GWT) and Prefrontal Cortex: Recent Developments. *Frontiers in Psychology*, *12*. https://www.frontiersin.org/articles/10.3389/fpsyg.2021.749868

Balduzzi, D., & Tononi, G. (2008). Integrated information in discrete dynamical systems: Motivation and theoretical framework. *PLoS Computational Biology*, *4*(6), e1000091. https://doi.org/10.1371/journal.pcbi.1000091

Baltieri, M., Biehl, M., Capucci, M., & Virgo, N. (2025). *A Bayesian Interpretation of the Internal Model Principle* (No. arXiv:2503.00511). arXiv. https://doi.org/10.48550/arXiv.2503.00511

Barrett, A. B., & Seth, A. K. (2011). Practical Measures of Integrated Information for Time-Series Data. *PLoS Computational Biology*, *7*(1), e1001052. https://doi.org/10.1371/journal.pcbi.1001052

Bellingrath, J. E. (2023). The Self-Simulational Theory of temporal extension. *Neuroscience of Consciousness*, *2023*(1), niad015. https://doi.org/10.1093/nc/niad015

Bergamaschi Ganapini, M., Campbell, M., Fabiano, F., Horesh, L., Lenchner, J., Loreggia, A., Mattei, N., Rossi, F., Srivastava, B., & Venable, K. B. (2025). Fast, slow, and metacognitive thinking in AI. *Npj Artificial Intelligence*, *1*(1), 27. https://doi.org/10.1038/s44387-025-00027-5

Bergson, H. (2001). *Time and Free Will: An Essay on the Immediate Data of Consciousness*. Dover Publications.

Bertschinger, N., Olbrich, E., Ay, N., & Jost, J. (2006). Information and closure in systems theory. In *Explorations in the Complexity of Possible Life. Proceedings of the 7th German Workshop of Artificial Life* (pp. 9–21). IOS Press.

Bialek, W., Nemenman, I., & Tishby, N. (2001). Predictability, complexity, and learning. *Neural Computation*, *13*(11), 2409–2463. https://doi.org/10.1162/089976601753195969

Biehl, M., & Virgo, N. (2023). Interpreting Systems as Solving POMDPs: A Step Towards a Formal Understanding of Agency. In C. L. Buckley, D. Cialfi, P. Lanillos, M. Ramstead, N. Sajid, H. Shimazaki, & T. Verbelen (Eds.), *Active*

Inference (pp. 16–31). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-28719-0_2

Bruner, J. (2003). Self-Making Narratives. In *Autobiographical Memory and the Construction of A Narrative Self*. Psychology Press.

Butlin, P., Long, R., Elmoznino, E., Bengio, Y., Birch, J., Constant, A., Deane, G., Fleming, S. M., Frith, C., Ji, X., Kanai, R., Klein, C., Lindsay, G., Michel, M., Mudrik, L., Peters, M. A. K., Schwitzgebel, E., Simon, J., & VanRullen, R. (2023). *Consciousness in Artificial Intelligence: Insights from the Science of Consciousness* (No. arXiv:2308.08708). arXiv. https://doi.org/10.48550/arXiv.2308.08708

Chalmers, D. J. (2024). *Could a Large Language Model be Conscious?* (No. arXiv:2303.07103). arXiv. https://doi.org/10.48550/arXiv.2303.07103

Chang, A. Y. C., Biehl, M., Yu, Y., & Kanai, R. (2020). Information Closure Theory of Consciousness. *Frontiers in Psychology*, *11*. https://www.frontiersin.org/articles/10.3389/fpsyg.2020.01504

Chen, S., Ma, S., Yu, S., Zhang, H., Zhao, S., & Lu, C. (2025). *Exploring Consciousness in LLMs: A Systematic Survey of Theories, Implementations, and Frontier Risks* (No. arXiv:2505.19806). arXiv. https://doi.org/10.48550/arXiv.2505.19806

Christoff, K., Irving, Z. C., Fox, K. C. R., Spreng, R. N., & Andrews-Hanna, J. R. (2016). Mind-wandering as spontaneous thought: A dynamic framework. *Nature Reviews. Neuroscience*, *17*(11), 718–731. https://doi.org/10.1038/nrn.2016.113

Conway, M. A., & Pleydell-Pearce, C. W. (2000). The construction of autobiographical memories in the self-memory system. *Psychological Review*, *107*(2), 261–288. https://doi.org/10.1037/0033-295X.107.2.261

Coull, J. T., Cheng, R.-K., & Meck, W. H. (2011). Neuroanatomical and Neurochemical Substrates of Timing. *Neuropsychopharmacology*, *36*(1), 3–25. https://doi.org/10.1038/npp.2010.113

Dehaene, S., Kerszberg, M., & Changeux, J.-P. (1998). A neuronal model of a global workspace in effortful cognitive tasks. *Proceedings of the National Academy of Sciences*, *95*(24), 14529–14534. https://doi.org/10.1073/pnas.95.24.14529

Dehaene, S., Lau, H., & Kouider, S. (2017). What is consciousness, and could machines have it? *Science*. https://doi.org/10.1126/science.aan8871

Fivush, R. (2011). The Development of Autobiographical Memory. *Annual Review of Psychology*, *62*(Volume 62, 2011), 559–582. https://doi.org/10.1146/annurev.psych.121208.131702

Fox, M. D., & Raichle, M. E. (2007). Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nature Reviews. Neuroscience*, *8*(9), 700–711. https://doi.org/10.1038/nrn2201

French, R. M. (1999). Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences*, *3*(4), 128–135. https://doi.org/10.1016/S1364-6613(99)01294-2

Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews. Neuroscience*, *11*(2), 127–138. https://doi.org/10.1038/nrn2787

Friston, K. (2018). Am I Self-Conscious? (Or Does Self-Organization Entail Self-Consciousness?). *Frontiers in Psychology*, *9*. https://doi.org/10.3389/fpsyg.2018.00579

Friston, K. J., Rosch, R., Parr, T., Price, C., & Bowman, H. (2018). Deep temporal models and active inference. *Neuroscience & Biobehavioral Reviews*, *90*, 486–501. https://doi.org/10.1016/j.neubiorev.2018.04.004

Friston, K., & Kiebel, S. (2009). Predictive coding under the free-energy principle. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *364*(1521), 1211–1221. https://doi.org/10.1098/rstb.2008.0300

Golden, R., Delanois, J. E., Sanda, P., & Bazhenov, M. (2022). Sleep prevents catastrophic forgetting in spiking neural networks by forming a joint synaptic weight representation. *PLOS Computational Biology*, *18*(11), e1010628. https://doi.org/10.1371/journal.pcbi.1010628

Hadsell, R., Rao, D., Rusu, A. A., & Pascanu, R. (2020). Embracing Change: Continual Learning in Deep Neural Networks. *Trends in Cognitive Sciences*, *24*(12), 1028–1040. https://doi.org/10.1016/j.tics.2020.09.004

Hayashi, M. J., van der Zwaag, W., Bueti, D., & Kanai, R. (2018). Representations of time in human frontoparietal cortex. *Communications Biology*, *1*(1), 233. https://doi.org/10.1038/s42003-018-0243-z

Husserl, E. (2011). *On the Phenomenology of the Consciousness of Internal Time* (J. B. Brough, Trans.). Springer.

James, W. (1890). *The Principles of Psychology*. Dover Publications.

Jolicoeur-Martineau, A. (2025). *Less is More: Recursive Reasoning with Tiny Networks* (No. arXiv:2510.04871). arXiv. https://doi.org/10.48550/arXiv.2510.04871

Juliani, A., Arulkumaran, K., Sasai, S., & Kanai, R. (2022). *On the link between conscious function and general intelligence in humans and machines* (No. arXiv:2204.05133). arXiv. http://arxiv.org/abs/2204.05133

Kent, L., & Wittmann, M. (2021). Time consciousness: The missing link in theories of consciousness. *Neuroscience of Consciousness*, *2021*(2), niab011. https://doi.org/10.1093/nc/niab011

Li, J., Fu, Y., Fan, L., Liu, J., Shu, Y., Qin, C., Yang, M., King, I., & Ying, R. (2025). *Implicit Reasoning in Large Language Models: A Comprehensive Survey* (No. arXiv:2509.02350). arXiv. https://doi.org/10.48550/arXiv.2509.02350

Li, Z.-Z., Zhang, D., Zhang, M.-L., Zhang, J., Liu, Z., Yao, Y., Xu, H., Zheng, J., Wang, P.-J., Chen, X., Zhang, Y., Yin, F., Dong, J., Li, Z., Bi, B.-L., Mei, L.-R., Fang, J., Liang, X., Guo, Z., … Liu, C.-L. (2025). *From System 1 to System 2: A Survey of Reasoning Large Language Models* (No. arXiv:2502.17419). arXiv. https://doi.org/10.48550/arXiv.2502.17419

Lingler, A., Talypova, D., Jokinen, J. P. P., Oulasvirta, A., & Wintersberger, P. (2024). Supporting Task Switching with Reinforcement Learning. *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 1–18. https://doi.org/10.1145/3613904.3642063

Mashour, G. A., Roelfsema, P., Changeux, J.-P., & Dehaene, S. (2020). Conscious Processing and the Global Neuronal Workspace Hypothesis. *Neuron*, *105*(5), 776–798. https://doi.org/10.1016/j.neuron.2020.01.026

McAdams, D. P. (1989). The Development of a Narrative Identity. In D. M. Buss & N. Cantor (Eds.), *Personality Psychology: Recent Trends and Emerging Directions* (pp. 160–174). Springer US. https://doi.org/10.1007/978-1-4684-0634-4_12

McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, *102*(3), 419–457. https://doi.org/10.1037/0033-295X.102.3.419

Mediano, P. A. M., Rosas, F. E., Bor, D., Seth, A. K., & Barrett, A. B. (2022). The strength of weak integrated information theory. *Trends in Cognitive Sciences*, *26*(8), 646–655. https://doi.org/10.1016/j.tics.2022.04.008

Millidge, B., Seth, A., & Buckley, C. L. (2022). *Predictive Coding: A Theoretical and Experimental Review* (No. arXiv:2107.12979). arXiv. https://doi.org/10.48550/arXiv.2107.12979

Naveed, H., Khan, A. U., Qiu, S., Saqib, M., Anwar, S., Usman, M., Akhtar, N., Barnes, N., & Mian, A. (2024). *A Comprehensive Overview of Large Language Models* (No. arXiv:2307.06435). arXiv. https://doi.org/10.48550/arXiv.2307.06435

Oizumi, M., Albantakis, L., & Tononi, G. (2014). From the Phenomenology to the Mechanisms of Consciousness: Integrated Information Theory 3.0. *PLOS*

Computational Biology, *10*(5), e1003588. https://doi.org/10.1371/journal.pcbi.1003588

Paton, J. J., & Buonomano, D. V. (2018). The Neural Basis of Timing: Distributed Mechanisms for Diverse Functions. *Neuron*, *98*(4), 687–705. https://doi.org/10.1016/j.neuron.2018.03.045

Protopapa, F., Hayashi, M. J., Kulashekhar, S., Zwaag, W. van der, Battistella, G., Murray, M. M., Kanai, R., & Bueti, D. (2019). Chronotopic maps in human supplementary motor area. *PLOS Biology*, *17*(3), e3000026. https://doi.org/10.1371/journal.pbio.3000026

Qi, B., Chen, X., Gao, J., Li, D., Liu, J., Wu, L., & Zhou, B. (2024). *Interactive Continual Learning: Fast and Slow Thinking* (No. arXiv:2403.02628). arXiv. https://doi.org/10.48550/arXiv.2403.02628

Rasch, B., & Born, J. (2013). About sleep's role in memory. *Physiological Reviews*, *93*(2), 681–766. https://doi.org/10.1152/physrev.00032.2012

Rosas, F., Boyd, A., & Baltieri, M. (2025, May 9). *AI in a vat: Fundamental limits of efficient world modelling for agent sandboxing and interpretability*. Reinforcement Learning Conference. https://openreview.net/forum?id=608omWjEcy&noteId=rz9t2OLX56

Rosas, F. E., Geiger, B. C., Luppi, A. I., Seth, A. K., Polani, D., Gastpar, M., & Mediano, P. A. M. (2024). *Software in the natural world: A computational approach to hierarchical emergence* (No. arXiv:2402.09090; Version 2). arXiv. https://doi.org/10.48550/arXiv.2402.09090

Shepherd, G. M. G., & Yamawaki, N. (2021). Untangling the cortico-thalamo-cortical loop: Cellular pieces of a knotty circuit puzzle. *Nature Reviews. Neuroscience*, *22*(7), 389–406. https://doi.org/10.1038/s41583-021-00459-3

Snell, C., Lee, J., Xu, K., & Kumar, A. (2024). *Scaling LLM Test-Time Compute Optimally can be More Effective than Scaling Model Parameters* (No. arXiv:2408.03314). arXiv. https://doi.org/10.48550/arXiv.2408.03314

Steriade, M., McCormick, D. A., & Sejnowski, T. J. (1993). Thalamocortical oscillations in the sleeping and aroused brain. *Science (New York, N.Y.)*, *262*(5134), 679–685. https://doi.org/10.1126/science.8235588

Sun, Y., Fujisawa, I., Juliani, A., Sakuma, J., & Kanai, R. (2024). *Remembering Transformer for Continual Learning* (No. arXiv:2404.07518). arXiv. https://doi.org/10.48550/arXiv.2404.07518

Tononi, G., & Cirelli, C. (2003). Sleep and synaptic homeostasis: A hypothesis. *Brain Research Bulletin*, *62*(2), 143–150.

https://doi.org/10.1016/j.brainresbull.2003.09.004
Tononi, G., & Cirelli, C. (2014). Sleep and the price of plasticity: From synaptic and cellular homeostasis to memory consolidation and integration. *Neuron*, *81*(1), 12–34. https://doi.org/10.1016/j.neuron.2013.12.025
VanRullen, R. (2016). Perceptual Cycles. *Trends in Cognitive Sciences*, *20*(10), 723–735. https://doi.org/10.1016/j.tics.2016.07.006
VanRullen, R., & Kanai, R. (2021). Deep learning and the Global Workspace Theory. *Trends in Neurosciences*, *44*(9), 692–704. https://doi.org/10.1016/j.tins.2021.04.005
VanRullen, R., & Koch, C. (2003). Is perception discrete or continuous? *Trends in Cognitive Sciences*, *7*(5), 207–213. https://doi.org/10.1016/S1364-6613(03)00095-0
VanRullen, R., Reddy, L., & Koch, C. (2006). The Continuous Wagon Wheel Illusion Is Associated with Changes in Electroencephalogram Power at ∼13 Hz. *Journal of Neuroscience*, *26*(2), 502–507. https://doi.org/10.1523/JNEUROSCI.4654-05.2006
Virgo, N., Biehl, M., Baltieri, M., & Capucci, M. (2025). *A “good regulator theorem” for embodied agents* (No. arXiv:2508.06326; Version 1). arXiv. https://doi.org/10.48550/arXiv.2508.06326
Wang, G., Li, J., Sun, Y., Chen, X., Liu, C., Wu, Y., Lu, M., Song, S., & Yadkori, Y. A. (2025). *Hierarchical Reasoning Model* (No. arXiv:2506.21734). arXiv. https://doi.org/10.48550/arXiv.2506.21734
Wang, L., Zhang, X., Su, H., & Zhu, J. (2024). A Comprehensive Survey of Continual Learning: Theory, Method and Application. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *46*(8), 5362–5383. https://doi.org/10.1109/TPAMI.2024.3367329
Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., Chi, E., Le, Q., & Zhou, D. (2023). *Chain-of-Thought Prompting Elicits Reasoning in Large Language Models* (No. arXiv:2201.11903). arXiv. https://doi.org/10.48550/arXiv.2201.11903
Whyte, C. J., Redinbaugh, M. J., Shine, J. M., & Saalmann, Y. B. (2024). Thalamic contributions to the state and contents of consciousness. *Neuron*, *112*(10), 1611–1625. https://doi.org/10.1016/j.neuron.2024.04.019
Wolff, A., Di Giovanni, D. A., Gómez-Pilar, J., Nakao, T., Huang, Z., Longtin, A., & Northoff, G. (2019). The temporal signature of self: Temporal measures of resting-state EEG predict self-consciousness. *Human Brain Mapping*, *40*(3),

789–803. https://doi.org/10.1002/hbm.24412

Zhang, Q., Lyu, F., Sun, Z., Wang, L., Zhang, W., Hua, W., Wu, H., Guo, Z., Wang, Y., Muennighoff, N., King, I., Liu, X., & Ma, C. (2025). *A Survey on Test-Time Scaling in Large Language Models: What, How, Where, and How Well?* (No. arXiv:2503.24235). arXiv. https://doi.org/10.48550/arXiv.2503.24235