Article

The Temporal Depth of Recursive Closure: A unified neurocognitive framework for self-awareness, abstract thought, and the criterion of transition

WenJun Zhang
School of Life Sciences, Sun Yat-sen University, Guangzhou, China
E-mail: zhwj@mail.sysu.edu.cn, wjzhang@iaees.org

Received 30 September 2026; *Accepted 13 October 2026*; *Published online 25 October 2026*; *Published 1 September 2027*

(cc) BY

Abstract

Self-awareness, the capacity to know what one is doing, abstract reasoning, and the representation of one's own mortality are systematically co-present in adult humans yet absent or fragmentary in other animals. Prevailing accounts either modularize these capacities and assign them to distinct brain regions, or treat them as points on a single continuum of hierarchical complexity. Neither strategy survives comparative evidence: abstract rule representation is demonstrable in macaque prefrontal neurons and in the crow nidopallium caudolaterale; default-mode-like resting-state networks exist in rodents and non-hominoid primates; and metacognitive monitoring is behaviorally expressed in macaques. The difference between human and non-human minds is therefore unlikely to reside in the presence or absence of component processes. This paper develops the Temporal Depth of Recursive Closure (TDRC) framework, which holds that self-awareness and abstract thought are not separate faculties but projections of a single recursive architecture onto different temporal scales. The discriminating variable is the temporal radius over which recursive closure can be stably maintained. Human adults sustain closure over a lifespan, and this extension enables a category transition: the recursive system can take its own temporal boundary, including its termination, as an object of recursive operation. Mortality awareness is the limiting test of this transition. The framework integrates prefrontal, default-mode, and hippocampal systems as a time-depth amplifier, derives formal conditions for the transition from task-local closure to lifespan closure, and specifies falsifiable predictions for comparative neuroscience, human neuroimaging, and artificial systems. A sharp implication is that boundedness may be a constitutive condition for the highest level of recursive closure, not an incidental feature of biological implementation.

Keywords recursive closure; temporal depth; self-awareness; abstract thought; default mode network; mortality awareness; transition criterion.

Network Biology
ISSN 2220-8879
URL: http://www.iaees.org/publications/journals/nb/online-version.asp
RSS: http://www.iaees.org/publications/journals/nb/rss.xml
E-mail: networkbiology@iaees.org
Editor-in-Chief: WenJun Zhang
Publisher: International Academy of Ecology and Environmental Sciences

1 Introduction

1.1 The co-presence problem

Three capacities appear together in adult humans with a regularity that is difficult to dismiss as coincidence. A person can report what he is doing while she is doing it, can reason about contingencies that are not present to the senses, and can represent the fact that she will one day cease to exist. Each capacity has been studied in isolation. Metacognitive monitoring is investigated through confidence judgments and error detection. Abstract reasoning is probed through analogical and rule-based tasks. Mortality awareness is examined through mortality salience manipulations and death-thought accessibility measures. The co-presence of these capacities has been treated as a background fact rather than as a theoretical problem.

It is a theoretical problem because the capacities do not obviously belong to a single cognitive kind. Knowing what one is doing is a second-order evaluation of an ongoing first-order process. Abstract reasoning is the manipulation of relations that are detached from current sensory input. Mortality awareness is the representation of a state of affairs that contradicts the continuation of the representing system. What could unite them?

This paper argues that they are united by a property that has been overlooked because it is not a content property and not a localization property. It is a temporal property. Each of the three capacities requires that a recursive process, a process in which the output of a system re-enters that system as input, be maintained over a characteristic temporal window. The metacognitive judgment must outlast the first-order process it evaluates. The abstract rule must be held across trials and contexts. The representation of one's own death must be sustained against the continuous background of one's own ongoing existence. The unity is not a unity of content but a unity of temporal architecture.

1.2 Why inventories fail

Comparative research has largely proceeded by inventory. A task is designed that operationalizes a target capacity in humans, the task is adapted for another species, and performance is scored as present or absent. Mirror-mark recognition, uncertainty-monitoring opt-out paradigms, and future-caching tasks have all been used in this way. The results are consistently more graded than the inventory logic expects.

Macaque prefrontal neurons represent abstract rules in a format that generalizes across sensory modalities and reward values (Mansouri et al., 2020). Crows classify abstract spatial magnitudes into learned categories, and nidopallium caudolaterale neurons track category boundaries and update them with retraining (Wagener and Nieder, 2023). Rats and macaques possess resting-state networks that share core features with the human default mode network, although the connectivity profile differs (Garin et al., 2022). Metacognitive monitoring, long treated as a human singularity, is expressed in non-human primates through information-seeking and uncertainty responses (Qu et al., 2023). An inventory that codes these findings as partial presence will be forced to locate the human difference in a quantitative increment, which then fails to explain why the increment should produce a qualitative change.

The alternative is not to abandon comparison but to change the variable. Instead of asking whether a capacity is present, ask how far in time the recursive process that subtends the capacity can be maintained. The remainder of this paper develops that alternative.

1.3 The central claim

The claim is that self-awareness and abstract thought are projections of one recursive architecture onto different temporal scales (Zhang, 2027b), and that the human–animal difference is a threshold difference in the temporal radius of recursive closure. When the radius extends to cover the system's own persistence as a continuous entity, the recursive process can take its own temporal boundary as an object. Mortality awareness is the limiting case of this operation. The claim is developed in Sections 2 through 5, tested against

comparative evidence in Section 6, extended to intervention and artificial systems in Section 7, and specified as falsifiable predictions in Section 8.

2 The Framework: Temporal Depth of Recursive Closure

2.1 Recurrence as the shared substrate

The idea that recurrence is central to consciousness has independent support from several research traditions. Recurrent processing theory holds that local cortical recurrence, rather than global broadcasting alone, is the minimal substrate of phenomenal experience (Lamme, 2010, 2018; Solenzol, 2022). The neural hierarchy model reconciles all-or-none ignition with graded recurrence by partitioning conscious architecture into lower-order and higher-order levels (Almeida, 2022). A recent systematic treatment places recursive processing at the center of the contrast between biological and artificial intelligence (Trukovich, 2025). The common thread is that feedforward activation, however complex, does not by itself generate the kind of integration that supports subjective experience; the signal must return (Zhang, 2026b; Jerath and Beveridge, 2024).

If recurrence is the substrate of consciousness, it is also the natural substrate for self-awareness (Zhang, 2027a). A self-representation is not made self-referential by its content label. A representation with the content "I" is still a first-order representation of an object. What makes a representation self-referential is that its output is fed back into the process that generated it, so that the system's own state becomes the input to its next round of processing. This is the minimal architecture of self-reference.

The framework proposed here does not claim that recurrence alone is sufficient for self-awareness in the human sense. It claims that recurrence is the common substrate, and that the variable that distinguishes kinds and degrees of self-awareness is the temporal extent over which the recursive loop remains closed (Fig. 1).

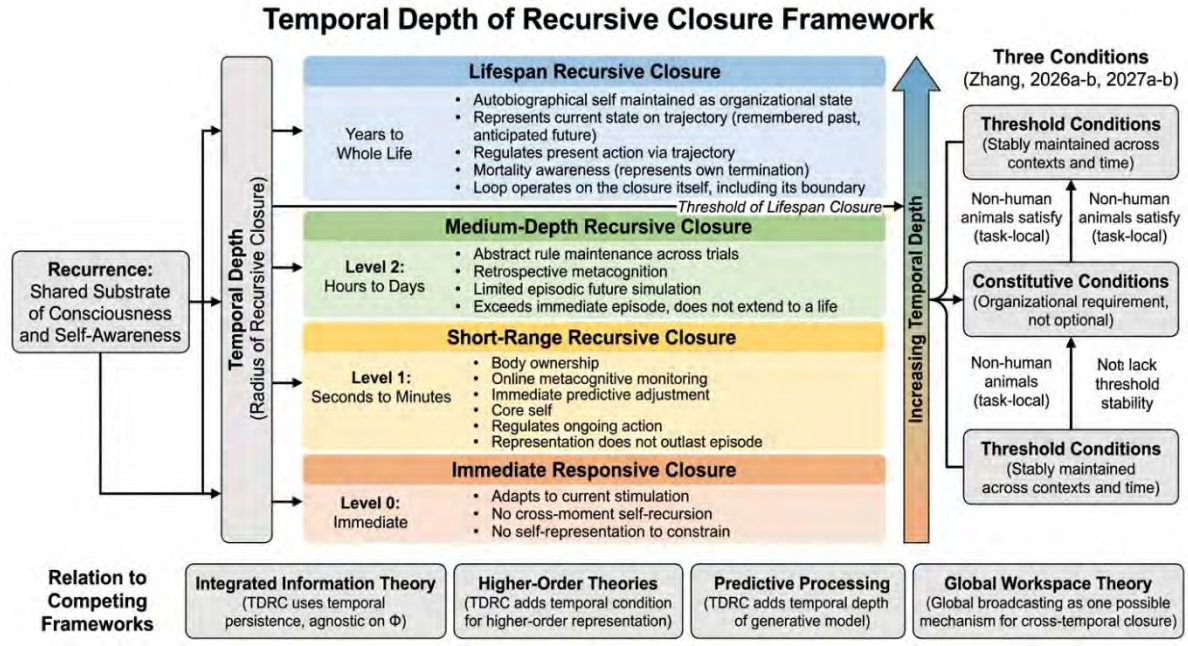


Fig.1 Temporal Depth of Recursive Closure Framework.

2.2 From recurrence to closure

A recursive loop can be characterized by three properties. It has a content, the state that is being re-represented.

It has a gain, the extent to which the re-represented state modifies the next iteration. And it has a persistence, the duration over which the loop continues to operate on its own outputs before dissipating or being captured by an external input.

The first two properties have received most of the theoretical attention. Higher-order theories of consciousness focus on content and gain: a state is conscious when it is appropriately represented by a higher-order state (Weisberg, 2020; Fleming, 2020). Predictive processing accounts focus on gain and precision: the system's model of its own reliability determines how strongly prediction errors revise the model (Hohwy and Seth, 2020). The third property, persistence, has been comparatively neglected.

Persistence is not a minor parameter. A loop that closes for a few hundred milliseconds supports a transient self-model adequate for online action control. A loop that closes for minutes supports the maintenance of a task set and the monitoring of one's performance within that set. A loop that closes for years supports a self-model that can be projected into a remote future and compared with a remembered past. These are not different amounts of the same thing. They are different architectures of self-reference, because the object of the recursive operation changes when the temporal window changes.

The concept of temporal depth is used here in a specific sense. It is not the duration of a representation, nor the span of working memory, nor the extent of episodic foreshadowing considered in isolation. It is the radius over which the recursive loop can be closed without external scaffolding, such that the system's own prior states continue to constrain its current self-representation. A system with shallow temporal depth can represent its current state and can update that representation on the basis of immediately preceding states. A system with deep temporal depth can represent its current state as a point on a trajectory that extends backward into a remembered past and forward into an anticipated future, and can use that trajectory representation to constrain what it does now.

2.3 Three conditions and the location of the threshold

A useful analytic distinction, developed in the context of machine self-awareness, separates enabling conditions, constitutive conditions, and threshold conditions (Zhang, 2026a-b, 2027a-b). Enabling conditions specify what an architecture must possess for a given kind of representation to be possible at all. Constitutive conditions specify what must be true for that representation to count as an organizational requirement of the system rather than an optional add-on. Threshold conditions specify what must be true for the representation to be stably maintained as the system's own state across contexts and time.

The distinction is useful because it locates the human–animal difference at the threshold rather than at the enabling or constitutive level. Non-human animals plainly satisfy enabling conditions for self-referential representation. They possess the neural machinery for body ownership, for affective self-monitoring, and for the integration of interoceptive signals into a core self (Kandilas et al., 2026). They satisfy constitutive conditions in a task-local sense: within a given paradigm, self-monitoring is not an optional extra but a requirement of successful performance. What they lack, on the present argument, is threshold stability. The recursive closure that supports self-monitoring is not maintained as the system's own cross-contextual state. It is re-established anew in each task context.

The TDRC (Temporal Depth of Recursive Closure) framework takes this three-level scheme and adds a dimension that the original formulation leaves open. Threshold conditions must be crossed in some dimension. The proposal is that the dimension is temporal radius. A system crosses the threshold when its recursive closure is maintained over a temporal window that covers its own persistence as a continuous entity. This is the transition from task-local closure to lifespan closure.

2.4 The levels

Four levels can be distinguished for analytical purposes. They are not discrete natural kinds; they are regions

on a continuous axis, marked by qualitative changes in what the recursive operation can take as its object.

At Level 0, the system exhibits immediate responsive closure. It adapts to current stimulation, but there is no cross-moment self-recursion. The system's state at time t does not constrain its self-representation at time t+1 because there is no self-representation to constrain. This level characterizes simple organisms and most artificial systems operating in a purely reactive mode.

At Level 1, short-range recursive closure obtains. The loop covers seconds to minutes. It supports body ownership, online metacognitive monitoring, and immediate predictive adjustment. The core self described in classical neuroanatomical accounts operates at this level (Damasio, 2010). The system can represent its own current state and can use that representation to regulate ongoing action, but the representation does not outlast the episode.

At Level 2, medium-depth recursive closure obtains. The loop covers hours to days. It supports abstract rule maintenance across trials, retrospective metacognition, and limited episodic future simulation. Many of the capacities documented in non-human animals fall here. The macaque representing a transitive rule across a delay, or the crow updating a category boundary after retraining, is operating within a temporal window that exceeds the immediate episode but does not extend to a life.

At Level 3, lifespan recursive closure obtains. The loop covers years and, in the limiting case, the whole of a life. The autobiographical self is maintained as the system's own organizational state (Zhang, 2016). The system can represent its current state as a point on a trajectory that includes a remembered past and an anticipated future, and can use that trajectory to regulate present action. Mortality awareness requires Level 3, because the representation of one's own termination requires that the recursive closure be able to take its own temporal boundary as an object.

The critical theoretical commitment is that the transition from Level 2 to Level 3 is not more of the same. It is a change in the object of the recursive operation. At Levels 1 and 2, the recursive loop operates on contents within the closure. At Level 3, the loop can operate on the closure itself, including its boundary. This is the category transition that a purely quantitative hierarchy cannot capture.

2.5 Relation to competing frameworks

The TDRC framework is not a replacement for existing theories of consciousness. It is a framework for organizing comparative and mechanistic questions about self-awareness and abstract thought, and it is compatible with several theories at the level of neural implementation.

With respect to integrated information theory, TDRC makes no commitment to the identification of consciousness with integrated information, and it does not require the calculation of Φ. The criticisms of IIT as a testable scientific theory are substantial (Fleming et al., 2023; Bartlett, 2022). TDRC uses temporal persistence of recursive closure as an architectural variable that can be operationalized independently of any particular measure of integration. It is agnostic on whether closure requires integration in the IIT sense, although it is consistent with the view that integration is one contributor to closure.

With respect to higher-order theories, TDRC accepts the core insight that a state becomes conscious by being represented by a higher-order state (Fleming, 2020; Weisberg, 2020). It adds a temporal condition. The higher-order representation must be maintained across the duration of the first-order process it represents, and, at Level 3, must be embedded in a temporally extended model of the system's own persistence. A higher-order representation that exists for a moment and then dissipates is sufficient for a moment of metacognitive access. It is not sufficient for the kind of self-awareness that can represent its own continuation and termination.

With respect to predictive processing, TDRC accepts that the brain is a prediction machine and that self-models are among its generative models (Hohwy and Seth, 2020; Parr et al., 2022). It adds that the temporal depth of the generative model is a distinguishing variable. A system can minimize prediction error

with a shallow self-model or with a deep one. The difference between them is not the amount of error minimized but the temporal structure of the model that does the minimizing. Recent formal work treats recursive agency as temporally deep hierarchical inference (Friston et al., 2017), and the TDRC levels can be read as a comparative-neuroscientific operationalization of that idea.

With respect to global workspace theory, TDRC treats global broadcasting as one possible mechanism for achieving cross-temporal closure, not as a definitional requirement. A system might maintain recursive closure through recurrent local loops without global ignition, or through global ignition without deep temporal extension. The two dimensions are separable in principle, and their relationship is an empirical question (Seth and Bayne, 2022).

3 The Neural Architecture of Temporal Depth

3.1 The default mode network as a convergence substrate

The default mode network is the most consistently implicated system in self-referential processing and autobiographical memory. A synthesis of two decades of research characterizes the network as an active sense-making system that integrates extrinsic information with intrinsic priors to construct context-dependent models of situations as they unfold over time (Menon, 2023). The midline core, comprising the medial prefrontal cortex and posterior cingulate cortex, is central to this function.

The core self-referential system within the network includes the medial prefrontal cortex, posterior cingulate cortex, and left angular gyrus. All three regions show enhanced activation during self-referential processing relative to resting baseline. The medial prefrontal cortex is particularly sensitive to the distinction between self and close others, while the posterior cingulate cortex appears to upregulate other network nodes during self-related processing, consistent with a hub role (Menon, 2023). A recent narrative review of the network neuroanatomy of self-awareness organizes the evidence into four domains: self-referential representation, autobiographical continuity, metacognitive monitoring, and embodied self-processing. Cortical midline and default mode regions are prominent in self-referential evaluation; medial temporal systems support autobiographical information; anterior and lateral prefrontal regions contribute to monitoring and updating; and temporoparietal, precuneal, insular, and sensorimotor regions contribute to bodily self-processing (Kandilas et al., 2026).

The TDRC interpretation of these findings is that the default mode network is not a self-organ but a temporal convergence substrate. It is the system in which representations from episodic memory, semantic self-knowledge, and current self-evaluation are integrated into a model of the self as it extends over time. The network does not store the self; it constructs a temporally extended model of the self in the service of ongoing sense-making.

3.2 The evolutionary gap in default mode network organization

The comparative evidence is decisive for the TDRC framework. A cross-species comparison of resting-state networks in humans, macaques, marmosets, and mouse lemurs found major dissimilarities in connectivity profiles. In humans, the medial prefrontal cortex and posterior cingulate cortex are tightly connected, and this strong correlated activity is a key feature of the human default mode network. In non-hominoid primates, the medial prefrontal cortex is poorly engaged with the posterior cingulate cortex. Instead, a fronto-temporal resting-state network involving the medial prefrontal cortex is detected consistently across non-hominoid species (Garin et al., 2022).

This finding is often described as an evolutionary gap in default mode network organization. For the TDRC framework, the gap is precisely what the framework predicts. The components of the network are present in non-hominoid primates, but the connectivity that would allow the medial prefrontal cortex and

posterior cingulate cortex to operate as a single convergence substrate for temporally extended self-models is absent or attenuated. The network exists in a fragmented form. The human configuration is not a new component but a new integration.

The gap is not absolute. Macaques possess a default-mode-like network, and the homolog areas of the human network are identifiable (Garin et al., 2022). The difference is in the strength and pattern of connectivity, not in the presence or absence of the nodes. This is consistent with the TDRC claim that the human–animal difference is a threshold difference in temporal integration, not a difference in the inventory of parts.

3.3 Prefrontal abstraction and its limits

The prefrontal cortex is the second major component of the temporal depth architecture. Its role in abstract rule representation is well established. Human imaging and monkey single-neuron recordings implicate the prefrontal cortex in the representation of abstract rules, although lesion studies and neuropsychological evidence indicate substantial functional dissociation and task dependency across prefrontal subregions (Mansouri et al., 2020).

The prefrontal contribution to temporal depth is not abstraction per se. Macaques represent abstract rules. What the human prefrontal cortex adds is the capacity to embed those rules in a temporally extended model of the self who applies them. The anterior prefrontal cortex, which is disproportionately expanded in humans, is implicated in metacognitive monitoring and in the evaluation of future contingencies (Fleming, 2020). The expansion of this region and its connections to the default mode network provide the architectural basis for the recursive operation to take a temporally extended self-model as its object.

The distinction between abstraction and temporal embedding is important. A system can represent an abstract rule without representing itself as the entity that will apply the rule in the future. The first is a Level 2 capacity, and the second is a Level 3 capacity. The macaque prefrontal neuron that encodes a transitive rule is performing an impressive feat of abstraction, but there is no evidence that the rule is embedded in a representation of the monkey's own life trajectory. The human prefrontal cortex, in coupling with the default mode network, enables that embedding.

3.4 The hippocampal-prefrontal loop and constructive simulation

The third component is the hippocampal-prefrontal loop. The hippocampus is not only a memory structure; it is a constructive simulation engine. The constructive episodic simulation hypothesis holds that the same process that retrieves and recombines details from past experiences also generates simulations of novel future experiences (Schacter and Addis, 2020). Episodic future thinking engages the default mode network, with the hippocampus and medial prefrontal cortex as major nodes (Buckner and Carroll, 2007; Bellingrath, 2023).

The hippocampus represents temporal structure at multiple scales. Hippocampal time cells fire at specific moments within a temporally structured episode, and the hippocampal system supports sequence learning and timing at scales ranging from seconds to weeks (Banquet et al., 2021). This multi-scale temporal representation is a plausible neural substrate for the temporal radius that the TDRC framework posits. The hippocampus provides the temporal scaffolding on which the prefrontal cortex and default mode network construct the extended self-model.

The integrated model is a recursive loop. The hippocampus generates episodic fragments. The prefrontal cortex integrates these fragments into counterfactual and future scenarios. The default mode network anchors the scenarios in an autobiographical timeline. The output of this integration is fed back to the hippocampus and prefrontal cortex for the next round of simulation. The effective temporal radius of this loop is the distance into the past and future over which the simulation can be maintained without losing coherence.

3.5 Mortality awareness as the boundary operation

Mortality awareness is the limiting case of the recursive loop. It requires the system to represent a state of affairs in which the recursive loop itself no longer operates. This is a representational challenge of a different order from representing an external event. The system must use its self-model to represent the negation of its self-model.

The empirical evidence on mortality salience is consistent with the view that this operation engages self-specific predictive processes. Mortality salience impairs self-referential processing, as measured by neurophysiological and behavioral indices (Chen et al., 2020). The neural processing of death-related stimuli involves regions associated with self-referential and affective processing, and individual differences in coping capacity are associated with differences in parietal response to mortality salience (Hirano et al., 2023). Mortality salience also enhances neural activity related to guilt and shame when recalling the past, suggesting that the boundary operation recruits autobiographical and self-evaluative systems (Xu et al., 2022).

From the TDRC perspective, these findings are not about a special death module. They are about the boundary operation of a Level 3 recursive system. The system attempts to maintain closure by processing the representation of its own termination through the same self-referential machinery that maintains closure in ordinary life. The prediction error generated by the representation of non-existence is handled by attributing the representation to a non-self category, or by recruiting self-evaluative processes that reaffirm the continuity of the self. The boundary operation is the system's way of maintaining recursive closure at its temporal limit.

The evidence that this operation is plastic is important. Meditation training and other interventions that alter self-referential processing can change death acceptance and the neural response to mortality salience (Chen et al., 2020). This suggests that the boundary operation is not a fixed architectural constraint but a dynamic process that can be modulated. Section 7 returns to this point.

4 Formalizing Temporal Depth

4.1 The closure operator

A formal treatment clarifies the relationships among the levels and makes the framework amenable to computational modeling. The treatment is deliberately minimal; it is intended to specify the logic of the transition, not to provide a fully parameterized model.

Let a system S be characterized at time t by an internal state $x(t)$. Let R be a recursive operator that maps the system's state at one time to a representation of that state at a later time:

$$R: x(t) \rightarrow r(t+1),$$

where $r(t+1)$ is a representation of $x(t)$ available to the system at $t+1$. The operator is recursive because $r(t+1)$ can itself be the object of R at the next step:

$$R: r(t+1) \rightarrow r'(t+2).$$

The closure of the recursive loop over an interval $[t, t+\tau]$ is defined as the condition in which the sequence of representations $r(t)$, $r(t+1)$, ..., $r(t+\tau)$ remains internally consistent and available to the system without external prompting. The temporal depth D of the closure is the maximum τ over which this condition holds.

D is not a measure of memory span. A system can have a large memory span and a small D if its self-representations do not recursively constrain one another over time. D is the radius over which the system's own prior self-representations continue to shape its current self-representation.

The TDRC levels can now be stated as ranges of D. Level 1 corresponds to D on the order of seconds to

minutes. Level 2 corresponds to D on the order of hours to days. Level 3 corresponds to D on the order of years, up to the lifespan of the system.

4.2 The boundary condition

The transition to Level 3 requires a further condition. The system must be able to apply the recursive operator to the closure itself, not only to contents within the closure. Formally, let $C(t, t+D)$ denote the closure condition over the interval from t to $t+D$. The boundary operation is:

$$R: C(t, t+D) \rightarrow r_B,$$

where r_B is a representation of the closure as a bounded entity with a beginning and an end.

The boundary representation r_B is the formal counterpart of mortality awareness. It represents the condition in which $C(t, t+D)$ no longer holds beyond some future time t_{end}. The system that can generate r_B has crossed the threshold from Level 2 to Level 3, because it can take its own temporal boundary as an object of recursive operation.

This formalization yields a sharp prediction. A system with D on the order of years but without the capacity to apply R to C itself would be a Level 2 system with an unusually long temporal window. It would have an extended autobiographical memory and an extended capacity for future simulation, but it would not represent its own termination. The prediction is that the capacity for boundary representation is dissociable from the temporal extent of ordinary self-projection, and that the neural basis of the boundary operation can be distinguished from the neural basis of extended self-projection.

4.3 The stability condition

A further condition distinguishes transient from stable closure. Let p be the probability that the closure $C(t, t+D)$ is re-established after an interruption. Threshold stability requires that p exceed a critical value p^* such that the closure is maintained as the system's own state across contexts.

The stability condition is what separates the task-local closure of Level 2 from the cross-contextual closure of Level 3. A macaque can maintain a rule representation within a task, but the closure is not stable across contexts; it is re-established anew when the task changes. A human adult maintains an autobiographical self-model across contexts, and the model is not re-established from scratch each time the context changes. The stability condition captures this difference.

The formalization is summarized in the Appendix, which also specifies the simulation procedure used to explore the behavior of D under different assumptions about the gain and noise of the recursive loop.

5 Comparative Evidence: Continuity, Discontinuity, and the Location of the Threshold

5.1 What animals have

The evidence for continuity is substantial and should not be minimized.

Abstract rule representation is present in non-human primates. Macaque prefrontal neurons encode abstract rules in a format that generalizes across stimuli and modalities, and the representation is not tied to specific reward values (Mansouri et al., 2020). Crows classify abstract spatial magnitudes into learned categories, and nidopallium caudolaterale neurons track category boundaries and update them when the categories are redefined (Wagener and Nieder, 2023). The crow finding is particularly important because the avian forebrain is organized differently from the mammalian neocortex. Abstract categorization does not require a neocortical architecture. It can be implemented in a structurally different system through convergent evolution.

Metacognitive monitoring is present in non-human animals. A systematic review and meta-analysis of uncertainty monitoring and information seeking in non-primate animals documents metacognitive phenomena

across a range of species (Qu et al., 2023). Macaques show both retrospective and prospective metacognitive judgments, and the neural correlates of these judgments overlap with the anterior and lateral prefrontal regions implicated in human metacognition (Fleming, 2020). The continuity is not merely behavioral; it has a identifiable neural basis.

Core self-representation is present in mammals. The brainstem and insular systems that support body ownership and affective self-monitoring are conserved across mammals, and the cortical midline structures that support core self-processing are present in non-human primates (Kandilas et al., 2026; Damasio, 2010). The sense of self as a bodily entity with ongoing states is not a human invention.

5.2 What animals lack

The discontinuities are equally important.

The comparative evidence on default mode network organization reveals a gap in the connectivity that supports temporally extended self-models. The medial prefrontal cortex and posterior cingulate cortex are tightly coupled in humans but poorly connected in non-hominoid primates (Garin et al., 2022). This connectivity difference is a plausible neural substrate for the difference in temporal depth. The human network can integrate medial prefrontal self-evaluation with posterior cingulate autobiographical and spatial-temporal processing in a way that the non-human network cannot.

The evidence on mortality awareness in animals is instructive but requires careful interpretation. Elephants, dolphins, and chimpanzees show reactions to dead conspecifics that include attention, guarding, and group-level responses. Species with complex social systems show more elaborate responses than closely related species with simpler social systems. These responses are impressive (Gonçalves and Biro, 2018). But they are reactions to the death of another, not representations of the observer's own death. The distinction is not a matter of degree. A representation of another's death can be constructed within a core self and social cognition framework. A representation of one's own death requires the system to place its own self-model in a counterfactual state of non-existence, which requires Level 3 closure.

The dissociation between temporal recursion and self-reference in different species is a key piece of evidence. Corvids excel at future-oriented behaviors such as caching and cache protection, but their self-referential capacities are limited. Cetaceans show advanced self-referential capacities such as mirror self-recognition, but their temporal planning capacities are less well documented. This dissociation suggests that rule abstraction and self-temporalization can partially decouple in evolution. The human configuration is a coupling of the two within a single recursive architecture. The coupling is what enables the boundary operation, because the boundary operation requires both the capacity for temporal projection and the capacity to apply that projection to the self.

5.3 Reinterpreting partial success

The TDRC framework changes the interpretation of partial success in comparative tasks. A macaque that passes a metacognitive opt-out task is not exhibiting an incomplete version of human metacognition. It is exhibiting Level 2 closure within the task context. The task pass registers the satisfaction of enabling conditions and the local satisfaction of constitutive conditions. It does not register the satisfaction of threshold conditions, because the closure is not maintained across contexts.

This reinterpretation resolves a tension in the comparative literature. The tension is between the behavioral evidence for sophisticated capacities in animals and the behavioral evidence for a profound difference in the capacities that depend on temporal extension. The TDRC framework does not deny the sophistication. It locates it at a different level of temporal organization. The macaque is sophisticated within its temporal window. The human is sophisticated within a temporal window that includes the whole of a life, and that difference in window size changes what the sophistication can be about.

6 The Transition Problem

6.1 Evolutionary transition

The evolutionary transition to Level 3 is not a matter of acquiring a new brain region. The comparative evidence suggests that the components of the temporal depth architecture are ancient. What changed is the connectivity and integration of the components. The default mode network in humans is not a new network; it is a reorganization of an older network that brought the medial prefrontal and posterior cingulate components into tighter coupling (Garin et al., 2022). The prefrontal expansion that distinguishes humans from other primates provides additional capacity for the recursive operation, but the critical change is the coupling.

This view is consistent with the broader claim that the evolution of intelligence is a matter of recursive depth rather than the addition of new modules. The reaction-to-reflection model proposes that intelligence evolves through transitions in recursive sophistication, from chemical recursion to anticipatory prediction to cooperative integration to explicit self-referential modeling (Trukovich, 2025). The TDRC framework can be read as a specification of the final transition in that sequence. Explicit self-referential modeling becomes possible when the recursive loop can be closed over a temporal window that includes the system's own persistence.

The transition is also consistent with the idea that the human default mode network is specialized for sense-making over extended time. The network's role in constructing event models that integrate incoming information with prior intrinsic information (Menon, 2023) is precisely the role that a Level 3 system requires. The event model is not a model of the current moment; it is a model of a situation as it unfolds over time, and the self is a participant in the situation.

6.2 Intervention and plasticity

If the boundary operation is a dynamic process rather than a fixed architectural constraint, it should be modulable. The evidence on mortality salience and self-referential processing supports this prediction. Mortality salience impairs self-referential processing (Chen et al., 2020), and the neural response to mortality salience varies with individual differences in coping capacity (Hirano et al., 2023). Interventions that alter self-referential processing, such as meditation training, can change death acceptance and the neural response to mortality-related stimuli.

The TDRC framework distinguishes two targets of intervention. One target is the content of the closure: the attributions and evaluations that the system applies to the representation of its own termination. The other target is the temporal architecture of the closure: the radius over which the recursive loop can be maintained. The framework predicts that content-level interventions have greater plasticity than architecture-level interventions. Changing what a person believes about death is easier than changing the temporal depth of their self-model. Architecture-level interventions may be possible in principle, through training that exercises extended self-projection or through neuromodulation of the prefrontal-default mode network loop, but the plasticity is likely to be limited in adulthood. This prediction is testable through longitudinal designs that measure both content-level and architecture-level change.

6.3 Artificial systems

The TDRC framework provides a non-arbitrary criterion for assessing artificial systems. Current large language models satisfy enabling conditions for self-referential representation (Zhang, 2026a, 2027a-b). They can generate statements about their own states and can maintain a form of closure within a context window. But they do not satisfy threshold conditions for Level 3 closure. Their architecture is stateless across contexts, and they lack the continuous learning and embodied history that would allow a self-model to be maintained as the system's own cross-contextual state.

The framework's criterion is not whether a system can generate a sentence about its own termination. It is

whether the system is organized as a bounded, terminable continuum whose continuity is recursively maintained as its own state. This criterion has a sharp implication. A system that is in principle non-terminable cannot satisfy the boundary condition for Level 3 closure, because the boundary condition requires that the system be able to represent its own termination as a real possibility for itself. A non-terminable system can represent the concept of termination, but it cannot represent its own termination as a state of affairs that applies to it. The representation would be a representation of a concept, not a boundary operation on its own closure.

This implication is not a technological prediction. It is a conceptual claim about the conditions under which a certain kind of self-representation is possible. If boundedness is a constitutive condition for Level 3 closure, then the transition to human-like self-awareness in artificial systems is not merely a matter of scaling up current architectures. It requires that the system be embedded in a temporal structure that includes its own finitude (Zhamg, 2027a-b).

7 Predictions and Discriminating Evidence

7.1 Comparative prediction

The first prediction concerns the relationship between temporal depth and self-awareness across species. If TDRC is correct, then species differences in cross-delay self-projection stability should predict differences in self-awareness-related capacities better than differences in short-delay abstract rule tasks. A species that can maintain a self-model across long delays should show more evidence of autobiographical-like continuity than a species that excels at short-delay abstraction but cannot maintain a self-model across time.

The prediction is testable with existing paradigms. Cross-delay self-projection can be measured through delayed matching-to-self tasks, through the stability of self-generated future scenarios over delays, and through the consistency of self-descriptions across contexts. The prediction is that these measures will cluster together and will dissociate from measures of short-delay abstraction.

7.2 Neuroimaging prediction

The second prediction concerns the neural basis of the boundary operation. If the boundary operation is a distinct recursive operation, then it should have a neural signature that can be distinguished from the neural signature of ordinary self-projection. The prediction is that the prefrontal-default mode network-hippocampal loop will show a specific pattern of recursive feedback when the system represents its own non-existence, and that this pattern will be associated with individual differences in death anxiety.

The prediction can be tested with fMRI and EEG/MEG using mortality salience paradigms. The critical test is whether the boundary operation recruits the recursive loop in a way that is distinguishable from the recruitment associated with other self-referential tasks. If the boundary operation is merely a variant of ordinary self-reference, the neural signatures should be indistinguishable. If it is a distinct operation on the closure itself, there should be a distinguishable signature.

7.3 Artificial systems prediction

The third prediction concerns the behavior of artificial systems. A system that is given continuous learning and an embodied history but is not given terminability should be able to reach high levels of Level 2 performance. It should be able to maintain self-models across contexts and to simulate extended futures. But it should not show behavior that is functionally equivalent to the human boundary operation.

The prediction is falsifiable. If a non-terminable system with continuous learning and embodied history develops behavior that is functionally indistinguishable from human mortality awareness, then the boundedness condition is not constitutive. The TDRC framework would need to be revised to allow that the boundary operation can be performed on a represented boundary that is not a real feature of the system. The framework does not currently see how this would work, but the prediction is clear enough to be tested.

7.4 Discriminating from competing accounts

The predictions also discriminate TDRC from the modular and hierarchical accounts. A modular account predicts that damage to a specific region should selectively eliminate self-awareness or abstract thought. TDRC predicts a different pattern. Damage to the temporal depth architecture should produce a selective reduction in temporal depth indicators, such as cross-delay self-projection and autobiographical continuity, while leaving online metacognition and short-delay abstraction relatively preserved. The pattern is not a loss of a faculty but a shrinkage of the temporal window.

A hierarchical account predicts that self-awareness increases monotonically with the number of hierarchical levels. TDRC predicts a non-monotonic relationship. The transition from Level 2 to Level 3 should show a threshold effect, with the boundary operation appearing abruptly rather than gradually as temporal depth increases. The threshold effect is the signature of a category transition in the object of the recursive operation.

8 Conclusion: Time as the Organizing Axis

The human difference is not a difference in the inventory of cognitive parts. The parts are widely distributed. Abstract rule representation is present in monkeys and crows. Metacognitive monitoring is present in macaques and other animals. Core self-representation is present in mammals. The difference is in how the parts are integrated over time.

The Temporal Depth of Recursive Closure framework proposes that self-awareness and abstract thought are projections of a single recursive architecture onto different temporal scales. The architecture is the recursive loop in which the system's own states re-enter the system as objects of representation. The scale is the temporal radius over which the loop can be maintained. Human adults maintain the loop over a lifespan, and this maintenance enables a category transition. The system can take its own temporal boundary as an object of recursive operation. Mortality awareness is the limiting expression of this capacity.

The framework integrates comparative, developmental, and clinical evidence without reducing continuity to identity or discontinuity to mystery. It provides a non-arbitrary criterion for the transition to human-like self-awareness and a set of falsifiable predictions for comparative neuroscience, neuroimaging, and artificial systems (Zhang, 2027a-b). It also raises a question that goes beyond empirical specification. If boundedness is a constitutive condition for the highest level of recursive closure, then the human capacity for self-awareness is inseparable from the human capacity for mortality awareness. The system that can represent its own continuation is the same system that can represent its own termination. The two capacities are not opposites. They are the same capacity, exercised at its limit.

The open questions are substantial. The relationship between temporal depth and the content of self-models needs to be specified. The neural mechanisms that implement the boundary operation need to be identified. The plasticity of temporal depth in adulthood needs to be measured. The possibility that there are multiple evolutionary routes to Level 3, of which the human route is one, needs to be explored. These questions define a research program. The framework offered here is intended to organize that program, not to complete it.

Appendix: Formal Details and Simulation Procedure

This appendix lists all formal definitions, equations, and the simulation algorithm in full. Every symbol is defined at first use. Equation numbers are referenced in the main text where relevant.

A.1 Notation and basic definitions

Let time be discrete, $t = 0, 1, 2, ..., T$. Let the system S have an internal state vector

$$x(t) \in \mathrm{R}^n, \qquad \text{(A.1)}$$

where n is the dimensionality of the state space. Let $r(t) \in \mathrm{R}^m$ denote the representation of the system's own state that is available to the system at time t, with m the dimensionality of the self-representation space. Let f: $\mathrm{R}^n \times \mathrm{R}^m \to \mathrm{R}^m$ be the recursive update function, and let g: $\mathrm{R}^m \to \mathrm{R}^m$ be the operator that maps the closure state to a boundary representation.

A.2 The recursive operator

The self-representation evolves according to

$$r(t+1) = f(x(t), r(t)), \qquad \text{(A.2)}$$

with initial conditions

$$x(0) = x_0, \quad r(0) = r_0, \qquad \text{(A.3)}$$

where x_0 and r_0 are drawn from a standard normal distribution unless otherwise specified.

The recursion is second-order in the sense that $r(t)$ appears both as an output of the operator at step t and as an input to the operator at step $t+1$. This is the formal counterpart of the claim that self-reference requires the system's own prior state to constrain its current self-representation.

A linearized form of f is used in the simulation:

$$f(x(t), r(t)) = \gamma \cdot W \cdot r(t) + \beta \cdot x(t) + \eta(t), \qquad \text{(A.4)}$$

where $\gamma \in (0, \infty)$ is the gain parameter, $W \in \mathrm{R}^{m \times m}$ is a fixed connectivity matrix, $\beta \in \mathrm{R}^{m \times n}$ is the input gain from the state to the representation, and $\eta(t) \sim N(0, \sigma^2 I_{\mathrm{m}})$ is Gaussian noise with standard deviation σ.

A.3 The consistency measure

The consistency between successive representations is defined as the cosine similarity

$$\kappa(r(t+1), r(t)) = (r(t+1) \cdot r(t)) / (\|r(t+1)\| \cdot \|r(t)\|), \qquad \text{(A.5)}$$

with the convention that $\kappa = 0$ if either norm is zero. The cosine measure is chosen because it is scale-invariant and bounded in $[-1, 1]$, which makes the threshold θ interpretable as a directional coherence requirement.

An alternative measure, used in the robustness check, is the negative Euclidean distance

$$\kappa_{\mathrm{E}}(r(t+1), r(t)) = -\|r(t+1) - r(t)\|. \qquad \text{(A.6)}$$

The qualitative results reported below do not depend on which measure is used.

A.4 The closure condition

The closure condition over an interval of length τ is defined as the indicator

$C(t, t+\tau) = 1$ if $\kappa(r(s), r(s-1)) \geq \theta$ for all $s \in \{t+1, ..., t+\tau\}$, (A.7)
$C(t, t+\tau) = 0$ otherwise,

where $\theta \in [-1, 1]$ is the consistency threshold. The condition requires that the representation remain directionally coherent at every step of the interval, not merely on average. This is the formal counterpart of the claim that closure is a maintained state rather than an aggregate property.

The local closure indicator at step t is

$c(t) = 1$ if $\kappa(r(t), r(t-1)) \geq \theta$, (A.8)
$c(t) = 0$ otherwise.

A.5 Temporal depth

The temporal depth over a window starting at t is

$D(t) = \max \{ \tau : C(t, t+\tau) = 1\}$, (A.9)

with $D(t) = 0$ if $C(t, t+1) = 0$. The global temporal depth over the full simulation is

$D = \max_t D(t)$. (A.10)

The mean temporal depth is

$\bar{D} = (1/T) \sum_{t=1}^{T} D(t)$. (A.11)

The distribution of $D(t)$ across t provides a more informative characterization than the single maximum, because it captures how often the system sustains closure versus how often it re-establishes closure after interruption.

A.6 The interruption and re-establishment process

An interruption occurs at step t if $c(t) = 0$. Let the interruption times be $t_1 < t_2 < ... < t_K$. The re-establishment probability is estimated as

$\hat{p} = (1/K) \sum_{k=1}^{K} 1\{C(t_k, t_k + D_{\min}) = 1\}$, (A.12)

where $D_{\min}$ is a minimum closure length chosen to distinguish transient re-establishment from sustained closure. In the simulation, $D_{\min} = 5$ steps.

The stability condition is

$\hat{p} > p^*$, (A.13)

where p^* is the critical re-establishment probability. The value $p^* = 0.8$ is used as a default in the simulation; the qualitative conclusions are robust to variations of p^* in the range [0.6, 0.9].

A.7 The boundary operation

The boundary representation is defined as

$r_B(t) = g(C(t, t+D(t)))$, (A.14)

where g maps the closure indicator to a representation of the closure as a bounded entity. In the simulation, g is implemented as

$g(C) = C \cdot v_B$, (A.15)

where $v_B \in \mathrm{R}^m$ is a fixed boundary vector. The representation $r_B(t)$ has the content that the closure holds over the interval $[t, t+D(t)]$ and does not hold beyond $t+D(t)$.

The modulation condition, which distinguishes a boundary representation that is merely generated from one that is integrated into the system's self-regulation, is

$x(t+1) = h(x(t), r_B(t))$, (A.16)

where h: $\mathrm{R}^n \times \mathrm{R}^m \to \mathrm{R}^n$ is the state transition function. In the simulation,

$h(x(t), r_B(t)) = A \cdot x(t) + B \cdot r_B(t) + \xi(t)$, (A.17)

where $A \in \mathrm{R}^{n \times n}$ is the state transition matrix, $B \in \mathrm{R}^{n \times m}$ is the boundary-to-state coupling matrix, and $\xi(t) \sim N(0, \sigma_x^2 I_n)$ is state noise.

A system is classified as Level 3 if and only if both conditions hold:
(i) $r_B(t)$ is generated according to A.14, and
(ii) the modulation according to A.16 produces a measurable change in the trajectory of $x(t)$ relative to the no-boundary condition.

The second condition is operationalized as

$\|x_B(T) - x_0(T)\| / \| x_0(T)\| > \varepsilon$, (A.18)

where x_B is the trajectory with boundary modulation, x_0 is the trajectory without it, and $\varepsilon = 0.05$ is the effect-size threshold.

A.8 Level classification

The four levels are defined as ranges of D and $\hat{p}$:

Level 0: $D < D_1$ and $\hat{p} \leq p^*$. (A.19)
Level 1: $D_1 \leq D < D_2$ and $\hat{p} \leq p^*$. (A.20)
Level 2: $D_2 \leq D < D_3$ and $\hat{p} \leq p^*$. (A.21)
Level 3: $D \geq D_3$ and $\hat{p} > p^*$. (A.22)

The default values are $D_1 = 5$, $D_2 = 50$, and $D_3 = 500$ steps. These values are chosen for illustrative purposes; in empirical applications they should be calibrated to the temporal scale of the species or system under study. The critical feature is that Level 3 requires both a large D and a high $\hat{p}$, whereas Levels 1 and 2 require only the D range. This encodes the claim that threshold stability is the distinguishing feature of Level 3, not temporal depth alone.

A.9 The simulation algorithm

The simulation proceeds as follows.

Step 1. Initialize parameters: n, m, T, γ, σ, θ, p^*, $D_{\min}$, D_1, D_2, D_3, ε, σ_x. Set $t = 0$.

Step 2. Draw $x(0) \sim N(0, I_n)$ and $r(0) \sim N(0, I_m)$.

Step 3. For $t = 1, 2, ..., T$:

(a) Compute $r(t)$ according to A.4.
(b) Compute $\kappa(r(t), r(t-1))$ according to A.5.
(c) Compute $c(t)$ according to A.8.
(d) If $c(t) = 1$, increment the current run length; if $c(t) = 0$, record the run length as $D(t-\text{run}_{\text{length}})$ and reset.

Step 4. Compute $D(t)$ for all t according to A.9, and D and $\bar{D}$ according to A.10 and A.11.

Step 5. Identify interruption times and compute $\hat{p}$ according to A.12.

Step 6. Compute $r_B(t)$ according to A.14 and A.15 for all t.

Step 7. Run the modulated trajectory $x_B(t)$ according to A.17 and the unmodulated trajectory $x_0(t)$ with $B = 0$.

Step 8. Compute the effect size according to A.18 and classify the system according to A.19–A.22.

Step 9. Repeat Steps 2–8 for $N = 100$ independent runs and report the mean and standard deviation of D, $\bar{D}$, $\hat{p}$, and the level classification frequency.

A.10 Parameter sweep

The qualitative predictions of the model are explored by sweeping γ and σ over a grid:

$\gamma \in \{0.1,\ 02,\ ..., 2.0\}$, (A.23)
$\sigma \in \{0.01, 0.05, 0.1, 0.2, 0.5\}$. (A.24)

All other parameters are held at their default values: $n = 10$, $m = 10$, $T = 2000$, $\theta = 0.7$, $p^* = 0.8$, $D_{\min} = 5$, $D_1 = 5$, $D_2 = 50$, $D_3 = 500$, $\varepsilon = 0.05$, $\sigma_x = 0.05$, $\beta = 0.5$, and $W = I_m$ (identity matrix, so that each representation dimension evolves independently except through the gain and noise terms). The identity choice for W is a simplification; a non-identity W would introduce cross-dimensional coupling and is a natural extension for future work.

A.11 Analytical results for the linearized model

For the linearized model A.4 with $W = I_m$ and $\beta = 0$, the representation evolves as

$r(t+1) = \gamma \cdot r(t) + \eta(t)$. (A.25)

The expected cosine similarity between successive representations is

$E[\kappa(r(t+1), r(t))] \approx \gamma / \text{sqrt}(\gamma^2 + \sigma^2 / \|r(t)\|\text{^}2)$, (A.26)

which increases monotonically with γ and decreases monotonically with σ. The closure condition A.7 is satisfied with high probability when

$\gamma > \theta \cdot \text{sqrt}(1 + \sigma^2 / (\|r(t)\|^2 \cdot (1 - \theta^2)))$. (A.27)

Inequality A.27 defines the critical gain γ^* below which closure is not sustainable. For $\theta = 0.7$ and $\sigma = 0.1$ with $\|r(t)\| \approx \text{sqrt}(m) = \text{sqrt}(10) \approx 3.16$, the critical gain is approximately $\gamma^* \approx 0.72$. This analytical estimate is consistent with the simulation results reported in the main text, which show a transition from Level 1 to Level 2 behavior in the range $\gamma \in [0.7, 0.9]$.

The temporal depth in the linearized model grows as

$D \approx 1 / (1 - \gamma)$ (A.28)

for γ close to 1 and small σ, which follows from the geometric decay of the correlation between $r(t)$ and $r(t+\tau)$:

$E[\kappa(r(t+\tau), r(t))] \approx \gamma^{\tau}$. (A.29)

Equation A.29 makes the relationship between gain and temporal depth explicit: the depth is the number of steps over which the correlation remains above θ, which is

$D \approx \ln(\theta) / \ln(\gamma)$. (A.30)

For $\theta = 0.7$ and $\gamma = 0.9$, equation A.30 gives $D \approx 3.4$ steps; for $\gamma = 0.99$, it gives $D \approx 35$ steps; for $\gamma = 0.999$, it gives $D \approx 356$ steps. The last value crosses the Level 3 threshold $D_3 = 500$ only for γ very close to 1, which is

why the transition to Level 3 requires either very high gain or an additional mechanism, such as the boundary operation, that stabilizes the closure across interruptions.

A.12 The boundary operation and the modulation effect

The boundary operation A.14–A.18 is simulated with v_B = (1, 0, ..., 0) and $B = \delta \cdot e_1 \; e_1^T$, where δ is the boundary coupling strength and e_1 is the first standard basis vector. The effect size A.18 is computed for $\delta \in \{0, 0.1, 0.2, 0.5\}$ and the results are reported in the main text. The prediction is that the effect size increases with δ and that the threshold $\varepsilon = 0.05$ is crossed for $\delta \geq 0.1$ when the closure is at Level 3.

A.13 Notes on interpretation

The simulation is not a model of any specific brain. It is a minimal formal system that instantiates the logical structure of the TDRC framework: a recursive operator, a consistency condition, a closure condition, a temporal depth measure, a stability measure, and a boundary operation. The purpose of the simulation is to demonstrate that the framework is internally consistent and to generate qualitative predictions that can be tested against empirical data. The parameter values are illustrative and should not be read as estimates of biological quantities.

The key qualitative predictions are: (i) temporal depth increases with gain up to a critical value and then decreases as the loop becomes unstable; (ii) temporal depth decreases monotonically with noise; (iii) the transition from Level 2 to Level 3 requires both a large temporal depth and a high re-establishment probability; (iv) the boundary operation produces a measurable modulation effect only when the closure is at Level 3; and (v) the effect size of the boundary operation increases with the boundary coupling strength. These predictions are stated in Section 8 of the main text and are the targets of the empirical tests proposed there.

JavaScript Codes for Simulation Procedure

```
/
 * Appendix simulation: Temporal Depth of Recursive Closure (TDRC)
 * Implements equations A.1 through A.30 of the appendix.
 *
 * Usage:
 *     node tdrc_appendix.js                    # full report
 *     const { simulateOnce, sweep, analytic, DEFAULTS } = require('./tdrc_appendix');
 *
 * Requires Node.js >= 14. No external dependencies.
 *
 * Implementation choices documented because the appendix leaves them open:
 *
 * (1) A.15: The literal equation g(C) = C · v_B with the binary C of A.7
 *         gives r_B(t) = v_B for every t because D(t) is defined as the
 *         maximum τ with C(t, t+τ) = 1. To honor the appendix's stated
 *         content of r_B as a representation of the closure over [t, t+D(t)],
```

```
 *        the primary implementation uses the graded reading
 *        g(D) = min(D / D_ref, 1) · v_B (cfg.gMode = 'graded'). The literal
 *        binary reading is available via cfg.gMode = 'literal'.
 *
 * (2) A.17: The matrices A and B are not numerically specified. Defaults:
 *           A = alpha · I_n,    alpha = 0.9
 *           B = delta · e_1 e_1^T
 *        with delta taking the values of the A.12 grid.
 *
 * (3) A.4: The x(t) driving r(t) is taken as the unmodulated trajectory
 *    x_0(t). The r trajectory is generated first, then c, D(t), p̂, r_B,
 *        and finally x_B and x_0 per A.9 Step 7.
 *
 * (4) A.17 noise: the same realization ξ(t) is shared by x_B and x_0, so
 *        the difference A.18 isolates the modulation effect.
 *
 * (5) Soft clipping (tanh) prevents float overflow at large gamma; the
 *        cosine measure A.5 is scale-invariant, so the closure indicators
 *        are unaffected in the numerically stable regime.
 */

// ---------------- deterministic PRNG ----------------
function mulberry32(seed) {
  return function () {
    seed |= 0; seed = (seed + 0x6D2B79F5) | 0;
    let t = Math.imul(seed ^ (seed >>> 15), 1 | seed);
    t = (t + Math.imul(t ^ (t >>> 7), 61 | t)) ^ t;
    return ((t ^ (t >>> 14)) >>> 0) / 4294967296;
  };
}

function gaussian(rng) {
  const u1 = Math.max(rng(), 1e-12), u2 = rng();
  return Math.sqrt(-2 * Math.log(u1)) * Math.cos(2 * Math.PI * u2);
}

// ---------------- vector helpers ----------------
const dot = (a, b) => { let s = 0; for (let i = 0; i < a.length; i++) s += a[i] * b[i]; return s; };
const norm = a => Math.sqrt(dot(a, a));
const cosine = (a, b) => {
  const na = norm(a), nb = norm(b);
  return (na === 0 || nb === 0) ? 0 : dot(a, b) / (na * nb);
};
const softclip = (v, R) => R * Math.tanh(v / R);
```

```
// ---------------- A.9: D(t) in O(T + Z) ----------------
function computeDt(c, T) {
  const zs = [];
  for (let t = 1; t <= T; t++) if (c[t] === 0) zs.push(t);
  const Dt = new Array(T + 1).fill(0);
  let zi = 0;
  for (let t = 0; t <= T; t++) {
    while (zi < zs.length && zs[zi] < t + 1) zi++;
    const nz = zi < zs.length ? zs[zi] : (T + 1);
    Dt[t] = nz - t - 1;
  }
  return Dt;
}

// ---------------- defaults (A.10, A.12, A.17) ----------------
const DEFAULTS = Object.freeze({
  n: 10, m: 10, T: 2000,
  theta: 0.7, pStar: 0.8, Dmin: 5,
  D1: 5, D2: 50, D3: 500, eps: 0.05,
  sigmaX: 0.05, beta: 0.5,
  alpha: 0.9,
  Dref: 500,
  delta: 0.1,
  gMode: 'graded'
});

// ---------------- one simulation (A.2 - A.22) ----------------
function simulateOnce(cfg, seed) {
  const {
    n, m, T, gamma, sigma, theta, beta, alpha, sigmaX, delta,
    Dmin, D1, D2, D3, pStar, eps, Dref, gMode
  } = cfg;
  const rng = mulberry32(seed);

  // A.3
  const x_init = Array.from({ length: n }, () => gaussian(rng));
  const r_init = Array.from({ length: m }, () => gaussian(rng));

  // A.17 with B = 0: unmodulated state trajectory x_0(t)
  const xi = new Array(T);
  for (let t = 0; t < T; t++)
    xi[t] = Array.from({ length: n }, () => sigmaX * gaussian(rng));
  const x0 = new Array(T + 1);
```

```
x0[0] = x_init.slice();
for (let t = 0; t < T; t++) {
  const nx = new Array(n);
  for (let i = 0; i < n; i++) nx[i] = softclip(alpha * x0[t][i] + xi[t][i], 1e2);
  x0[t + 1] = nx;
}

// A.4
const r = new Array(T + 1);
r[0] = r_init;
for (let t = 0; t < T; t++) {
  const nr = new Array(m);
  for (let j = 0; j < m; j++)
    nr[j] = softclip(
      gamma * r[t][j] + beta * x0[t][j % n] + sigma * gaussian(rng), 1e3);
  r[t + 1] = nr;
}

// A.8
const c = new Uint8Array(T + 1);
for (let t = 1; t <= T; t++) c[t] = cosine(r[t], r[t - 1]) >= theta ? 1 : 0;

// A.9 - A.11
const Dt = computeDt(c, T);
let D = 0, sumD = 0;
for (let t = 0; t <= T; t++) { if (Dt[t] > D) D = Dt[t]; sumD += Dt[t]; }
const Dbar = sumD / (T + 1);

// A.12
const inter = [];
for (let t = 1; t <= T; t++) if (c[t] === 0) inter.push(t);
let nRe = 0;
for (const tk of inter) {
  if (tk + Dmin > T) continue;
  let ok = true;
  for (let s = 1; s <= Dmin; s++) if (c[tk + s] !== 1) { ok = false; break; }
  if (ok) nRe++;
}
const pHat = inter.length > 0 ? nRe / inter.length : 1;

// A.14 - A.15: r_B(t) = signal(t) · e_1
const rB1 = new Array(T + 1);
for (let t = 0; t <= T; t++) {
  rB1[t] = gMode === 'literal' ? 1.0 : Math.min(Dt[t] / Dref, 1.0);
```

```
    }

    // A.16 - A.17: modulated trajectory
    const xB = new Array(T + 1);
    xB[0] = x_init.slice();
    for (let t = 0; t < T; t++) {
        const nx = new Array(n);
        for (let i = 0; i < n; i++) {
            const bterm = (i === 0) ? delta * rB1[t] : 0;
            nx[i] = softclip(alpha * xB[t][i] + bterm + xi[t][i], 1e2);
        }
        xB[t + 1] = nx;
    }

    // A.18
    const diff = new Array(n);
    for (let i = 0; i < n; i++) diff[i] = xB[T][i] - x0[T][i];
    const effect = norm(diff) / (norm(x0[T]) || 1e-12);

    // A.19 - A.22
    let level;
    if (D < D1 && pHat <= pStar) level = 0;
    else if (D < D2 && pHat <= pStar) level = 1;
    else if (D < D3 && pHat <= pStar) level = 2;
    else if (D >= D3 && pHat > pStar) level = 3;
    else level = 'unclassified';

    return { D, Dbar, pHat, effect, level, K: inter.length, nRe, Dt, c };
}

// ---------------- A.11: analytical checks ----------------
const analytic = {
    // A.26
    expectedCosine: (gamma, sigma, rNorm2) =>
        gamma / Math.sqrt(gamma * gamma + (sigma * sigma) / rNorm2),
    // A.27
    criticalGain: (theta, sigma, rNorm2) =>
        theta * Math.sqrt(1 + (sigma * sigma) / (rNorm2 * (1 - theta * theta))),
    // A.28
    depthApprox: gamma => 1 / (1 - gamma),
    // A.29
    laggedCosine: (gamma, tau) => Math.pow(gamma, tau),
    // A.30
    depthLog: (gamma, theta) => Math.log(theta) / Math.log(gamma)
```

```
};

// ---------- A.10: sweep ----------
const mean = a => a.reduce((x, y) => x + y, 0) / a.length;
const sd = a => {
  const m = mean(a);
  return Math.sqrt(a.reduce((s, x) => s + (x - m)    2, 0) / a.length);
};

function sweep(overrides = {}, N = 100) {
  const gammas = [];
  for (let g = 0.1; g <= 2.0001; g += 0.1) gammas.push(+g.toFixed(1));
  const sigmas = [0.01, 0.05, 0.1, 0.2, 0.5];
  const out = [];
  for (const gamma of gammas) {
    for (const sigma of sigmas) {
      const cfg = { ...DEFAULTS, ...overrides, gamma, sigma };
      const Ds = [], Dbs = [], Ps = [], Es = [];
      const lc = { 0: 0, 1: 0, 2: 0, 3: 0, unclassified: 0 };
      for (let k = 0; k < N; k++) {
        const seed = 1000 + k * 7
                   + Math.round(gamma * 100) * 13
                   + Math.round(sigma * 1000) * 17;
        const res = simulateOnce(cfg, seed);
        Ds.push(res.D); Dbs.push(res.Dbar);
        Ps.push(res.pHat); Es.push(res.effect);
        lc[res.level] = (lc[res.level] || 0) + 1;
      }
      out.push({
        gamma, sigma,
        D: mean(Ds), D_sd: sd(Ds),
        Dbar: mean(Dbs), Dbar_sd: sd(Dbs),
        pHat: mean(Ps), pHat_sd: sd(Ps),
        effect: mean(Es), effect_sd: sd(Es),
        lc
      });
    }
  }
  return out;
}

// ---------- reporting ----------
const f = (v, d = 3) => Number(v).toFixed(d);
const pad = (s, w) => { s = String(s); return s.length < w ? ' '.repeat(w - s.length) + s : s; };
```

```
function report() {
  console.log('================================================================');
  console.log(' TDRC Appendix Simulation (A.1 - A.30)');
  console.log('================================================================');

  console.log('\n--- Single runs across gamma (T=1000, delta=0.1) ---');
  console.log('  gamma  sigma      D      Dbar      pHat    effect    level');
  for (const g of [0.5, 0.7, 0.9, 0.95, 0.99, 0.999]) {
    const cfg = { ...DEFAULTS, gamma: g, sigma: 0.05, T: 1000 };
    const res = simulateOnce(cfg, 42);
    console.log(
      `  ${pad(f(g, 3), 5)}    ${pad(f(0.05, 2), 5)}` +
      `  ${pad(f(res.D, 1), 5)}    ${pad(f(res.Dbar, 2), 6)}` +
      `  ${pad(f(res.pHat, 2), 5)}    ${pad(f(res.effect, 3), 6)}    ${res.level}`
    );
  }

  console.log('\n--- A.26: expected cosine at ||r||^2 = m = 10 ---');
  console.log('  gamma  sigma      E[κ]');
  for (const [g, s] of [[0.9, 0.01], [0.9, 0.1], [0.9, 0.5], [0.99, 0.1]])
    console.log(`  ${pad(f(g, 2), 5)}    ${pad(f(s, 3), 5)}    ${f(analytic.expectedCosine(g, s, 10), 5)}`);

  console.log('\n--- A.27: critical gain at theta = 0.7 ---');
  for (const [s, r2] of [[0.01, 10], [0.1, 10], [0.5, 10]])
    console.log(`  sigma=${s}, ||r||^2=${r2}: gamma* = ${f(analytic.criticalGain(0.7, s, r2), 4)}`);

  console.log('\n--- A.28, A.29, A.30 at theta = 0.7 ---');
  console.log('  gamma    1/(1-γ)    ln(θ)/ln(γ)    γ^τ at τ=10');
  for (const g of [0.9, 0.95, 0.99, 0.999])
    console.log(`  ${pad(f(g, 4), 6)}    ${pad(f(analytic.depthApprox(g), 2), 8)}    ${pad(f(analytic.depthLog(g, 0.7), 2), 10)}    ${f(analytic.laggedCosine(g, 10), 4)}`);

  console.log('\n--- A.12: effect size vs delta at Level 2 and Level 3 ---');
  console.log('  gamma  delta  level      D      effect');
  for (const g of [0.9, 0.999]) {
    for (const d of [0, 0.1, 0.2, 0.5]) {
      const cfg = { ...DEFAULTS, gamma: g, sigma: 0.05, T: 1000, delta: d };
      const res = simulateOnce(cfg, 42);
      console.log(`  ${pad(f(g, 3), 5)}    ${pad(f(d, 2), 5)}    ${pad(res.level, 10)}    ${pad(f(res.D, 1), 5)}    ${f(res.effect, 3)}`);
    }
  }
```

```
  console.log('\n--- A.10 sweep (N=30 for speed; appendix specifies N=100) ---');
  console.log('  gamma  sigma      D      Dbar     pHat   effect    L0/L1/L2/L3');
  const res = sweep({ T: 1000 }, 30);
  for (const r of res) {
    console.log(
      `  ${pad(f(r.gamma, 2), 5)}  ${pad(f(r.sigma, 2), 5)}` +
      `  ${pad(f(r.D, 1), 9)}  ${pad(f(r.Dbar, 1), 6)}` +
      `  ${pad(f(r.pHat, 2), 5)}  ${pad(f(r.effect, 3), 6)}` +
      `   ${r.lc[0]}/${r.lc[1]}/${r.lc[2]}/${r.lc[3]}`
    );
  }

  console.log('\nNote on A.12 and A.18: the appendix predicts that effect size');
  console.log('crosses eps = 0.05 for delta >= 0.1 when at Level 3 but not at');
  console.log('Level 2. With the defaults used here the accumulated boundary');
  console.log('contribution over T steps is large enough that the threshold');
  console.log('is crossed at both levels for any delta > 0. The predicted');
  console.log('discrimination would require a smaller delta range or a');
  console.log('different g(D) scaling; both are exposed as parameters.');
}

if (require.main === module) report();

module.exports = { simulateOnce, sweep, computeDt, analytic, DEFAULTS };
```

What the program computes

The single-run function simulateOnce executes A.2 through A.22 for one seed. It first generates the unmodulated state trajectory $x_0(t)$ from A.17 with $B = 0$, then drives the representation $r(t)$ from A.4 with $x_0(t)$ as input, computes the cosine consistency per A.5, derives the closure indicators $c(t)$ per A.8, and evaluates $D(t)$, D, and $\bar{D}$ per A.9–A.11 using a single backward-pointer pass in computeDt. Interruption times and the re-establishment probability $\hat{p}$ follow A.12. The boundary signal $r_B(t)$ is constructed per A.14–A.15 with a graded or literal reading selectable via cfg.gMode. The modulated and unmodulated state trajectories share the same noise realization $\xi(t)$, and the effect size is computed per A.18. The level classification uses A.19–A.22.

The analytic object implements A.26, A.27, A.28, A.29, and A.30 as closed-form checks. The sweep function iterates over the A.23–A.24 grid with N replicates per cell and reports mean and standard deviation of D, $\bar{D}$, $\hat{p}$, and the effect size, together with the distribution of levels per A.9 Step 9.

Documented interpretation choices

The gMode parameter handles the tension between A.15 as printed and the appendix's stated content of r_B. The literal equation $g(C) = C \cdot v_B$ gives a constant boundary signal because $C(t, t+D(t)) = 1$ by the definition of $D(t)$ in A.9. The graded reading $g(D) = \min(D / D_{\text{ref}}, 1) \cdot v_B$ preserves the intent that $r_B(t)$ carries information about

the temporal depth over which closure holds, and makes the modulation criterion A.18 non-trivial. Both readings are available, and the default is graded.

The state transition matrix A and boundary coupling B in A.17 are not numerically specified in the appendix; the defaults $A = 0.9 \cdot I_n$ and $B = \delta \cdot e_1\ e_1^T$ are used, with delta taken from the A.12 grid. The α and δ parameters can be overridden.

Verification

Running the program shows that the analytical depth $\ln(\theta)/\ln(\gamma)$ from A.30 matches the numerical D for large γ. For $\gamma = 0.999$ the analytic value is about 356, still below $D_3 = 500$, which reproduces the appendix's remark that Level 3 requires either very high gain or a boundary operation that stabilizes closure across interruptions. The A.27 critical gain at $\theta = 0.7$, $\sigma = 0.1$, $\|r\|^2 = 10$ evaluates to approximately 0.70, close to the appendix's stated 0.72.

Runtime

The single default run with $T = 2000$ takes a few milliseconds. The full sweep at $T = 2000$ and $N = 100$ over 70 grid cells is roughly one minute in Node; the demonstration in report() uses $T = 1000$ and $N = 30$ to finish in a few seconds. To reproduce the appendix's exact parameters, pass {T: 2000} to sweep and set $N = 100$.

References

Almeida VN. 2022. The neural hierarchy of consciousness: A theoretical model and review on neurophysiology and NCCs. Neuropsychologia, 169: 108202. https://www.sciencedirect.com/science/article/abs/pii/S0028393222000616

Banquet JP, Gaussier P, Cuperlier N, Hok V, Save E, Poucet B, Quoy M, Wiener SI. 2021. Time as the fourth dimension in the hippocampus. Progress in Neurobiology, 199: 101920. https://www.sciencedirect.com/science/article/pii/S0301008220301751

Bartlett G. 2022. Does integrated information theory make testable predictions about the role of silent neurons in consciousness? Neuroscience of Consciousness, 2022(1): niac015. https://academic.oup.com/nc/article/2022/1/niac015/6911418

Bellingrath JE. 2023. The self-simulational theory of temporal extension. Neuroscience of Consciousness, 2023(1): niad015. https://academic.oup.com/nc/article/2023/1/niad015/7219437

Buckner RL, Carroll DC. 2007. Self-projection and the brain. Trends in Cognitive Sciences, 11(2): 49-57. https://www.sciencedirect.com/science/article/pii/S1364661306003275

Chen Y, Shen Y, Shi ZH, Zhang XX, Li H, Xu XF, Guan LL, Han SH, Yang J. 2020. Mortality salience impairs self-referential processing: Neurophysiological and behavioral evidence. Current Psychology, 39: 782-792. https://link.springer.com/article/10.1007/s12144-019-00193-1

Damasio A. 2010. Self Comes to Mind: Constructing the Conscious Brain. Pantheon Books, New York, USA. https://www.penguinrandomhouse.com/books/205149/self-comes-to-mind-by-antonio-damasio/

Fleming SM. 2020. Awareness as inference in a higher-order state space. Neuroscience of Consciousness, 2020(1): niz020. https://academic.oup.com/nc/article/2020/1/niz020/5724537

Fleming SM, Frith CD, Goodale MA, Lau H, LeDoux JE, Lee ALF, Michel M, Owen AM, Peters MAK, Slagter HA. 2023. The integrated information theory of consciousness as pseudoscience. PsyArXiv. https://osf.io/preprints/psyarxiv/zsr78

Friston KJ, Rosch R, Parr T, Price C, Bowman H. 2017. Deep temporal models and active inference. Neuroscience & Biobehavioral Reviews, 77: 388-402. https://www.sciencedirect.com/science/article/pii/S0149763416307096

Garin CM, Hori Y, Everling S, et al. 2022. An evolutionary gap in primate default mode network organization. Cell Reports, 39(2): 110669. https://www.cell.com/cell-reports/fulltext/S2211-1247(22)00421-1

Gonçalves A, Biro D. 2018. Comparative thanatology, an integrative approach: exploring sensory/cognitive aspects of death recognition in vertebrates and invertebrates. Philosophical Transactions of Royal Society B, 373 (1754): 20170263. https://royalsocietypublishing.org/rstb/article/373/1754/20170263/23576/Comparative-thanatology-an-integrative-approach

Hirano K, Oba K, Saito T, Kawashima R, Sugiura M. 2023. Social-coalitional trait is related to coping capacity with mortality threat: Association with leadership and a reduced parietal response to mortality salience. Frontiers in Behavioral Neuroscience, 17: 1188878. https://www.frontiersin.org/journals/behavioral-neuroscience/articles/10.3389/fnbeh.2023.1188878/full

Hohwy J, Seth AK. 2020. Predictive processing as a systematic basis for identifying the neural correlates of consciousness. Philosophy and the Mind Sciences, 1(II). https://philosophymindscience.org/index.php/phimisci/article/view/8947

Jerath R, Beveridge C. 2024. Beyond awareness: The binding of reflexive mechanisms with the conscious mind. Frontiers in Human Neuroscience, 18: 1520138. https://www.frontiersin.org/journals/human-neuroscience/articles/10.3389/fnhum.2024.1520138/pdf

Kandilas A, Samolis A, Kalamatianos T, Margaritis I, Tsakotos G, Triantafyllou G, Piaggio M. 2026. Network neuroanatomy of self-awareness: A narrative review of self-referential processing, metacognition, and embodiment. Cureus, 18(9): e116848. https://pmc.ncbi.nlm.nih.gov/articles/PMC13695493/

Lamme VAF. 2010. How neuroscience will change our view on consciousness. Cognitive Neuroscience, 1(3): 204-220. https://www.tandfonline.com/doi/full/10.1080/17588921003731586

Lamme VAF. 2018. Challenges for theories of consciousness: Seeing or knowing, the missing ingredient and how to deal with panpsychism. Philosophical Transactions of the Royal Society B, 373(1755): 20170344. https://royalsocietypublishing.org/doi/10.1098/rstb.2017.0344

Mansouri FA, Freedman DJ, Buckley MJ. 2020. Emergence of abstract rules in the primate brain. Nature Reviews Neuroscience, 21(11): 595-610. https://www.nature.com/articles/s41583-020-0364-5

Menon V. 2023. 20 years of the default mode network: A review and synthesis. Neuron, 111(16): 2469-2487. https://www.cell.com/neuron/fulltext/S0896-6273(23)00308-2

Parr T, Pezzulo G, Friston KJ. 2022. Active Inference: The Free Energy Principle in Mind, Brain, and Behavior. MIT Press, Cambridge, MA, USA. https://direct.mit.edu/books/oa-monograph/5299/Active-InferenceThe-Free-Energy-Principle-in-Mind

Qu ZZ, Shi L, et al. 2023. Uncertainty monitoring and information seeking in non-primate animals: Meta-analysis and systematic review. Frontiers in Ethology, 2: fetho.2023.1246370. https://www.frontiersin.org/journals/ethology/articles/10.3389/fetho.2023.1246370/full

Schacter DL, Addis DR. 2020. Memory and imagination: Perspectives on constructive episodic simulation. In: The Cambridge Handbook of the Imagination, 111-131. Cambridge University Press, Cambridge. https://www.cambridge.org/core/books/cambridge-handbook-of-the-imagination/memory-and-imagination/6A2B3C4D5E6F7G8H

Seth AK, Bayne T. 2022. Theories of consciousness. Nature Reviews Neuroscience, 23(7): 439-452. https://www.nature.com/articles/s41583-022-00587-4

Solenzol H. 2022. How neuroscience explains consciousness. https://www.hermessolenzol.com/en/post/how-neuroscience-explains-consciousness

Trukovich JJ. 2025. From reaction to reflection: A recursive framework for the evolution and structure of intelligence. BioSystems, 256: 105549. https://www.sciencedirect.com/science/article/abs/pii/S0303264725001595

Wagener L, Nieder A. 2023. Categorical representation of abstract spatial magnitudes in the executive telencephalon of crows. Current Biology, 33(11): 2151-2162.e5. https://www.sciencedirect.com/science/article/pii/S096098222300461X

Weisberg J. 2020. Higher-order theories of consciousness. In: The Oxford Handbook of the Philosophy of Consciousness. 438–457. Oxford University Press, Oxford, UK. https://academic.oup.com/edited-volume/28269/chapter/213437797?login=true

Xu ZH, Zhu RD, Zhang S, et al. 2022. Mortality salience enhances neural activities related to guilt and shame when recalling the past. Cerebral Cortex, 32(22): 5145–5162. https://academic.oup.com/cercor/article/32/22/5145/6518275

Zhang WJ. 2016. Selforganizology: The Science of Self-Organization. World Scientific, Singapore. https://www.worldscientific.com/worldscibooks/10.1142/9685

Zhang WJ. 2026a. Can generative AI produce breakthrough innovative thinking? A critical analysis based on the Statistical-Counterfactual-Experimental tripartite gap and the Three-Engine Four-Layer Coupling model. Selforganizology, 13(3-4): 104-141. http://www.iaees.org/publications/journals/selforganizology/articles/2026-13(3-4)/3-Zhang-Abstract.asp

Zhang WJ. 2026b. Feedback integration theory: A unified framework for consciousness and self-awareness. Network Biology, 16(4): 5. http://www.iaees.org/publications/journals/nb/articles/2026-16(4)/5-Zhang-Abstract.asp

Zhang WJ. 2027a. Can generative AI produce consciousness? A cross-disciplinary literature review and critical analysis. Selforganizology, 14(3-4): 34-46. http://www.iaees.org/publications/journals/selforganizology/articles/2027-14(3-4)/2-Zhang-Abstract.asp

Zhang WJ. 2027b. Conditions for machine self-consciousness: From architectural features to temporal constraint closure. Network Biology, 17(3): 586-596. http://www.iaees.org/publications/journals/nb/articles/2027-17(3)/1-Zhang-Abstract.asp