Intelligence Without Consciousness: A Diagnostic Paper on LLMs, Amoebae, and the Attractor Framework [F] (2026)

Robert Galida — June 2026

Abstract

The attractor framework defines intelligence as the ability to navigate a constraint field — to update behavior in response to perturbations and find persistent trajectories. Consciousness, within this framework, requires additional properties: a unified dissipative body, a persistent self-model, phenomenal valence (subjective liking/disliking), and subjective experience. This paper applies that diagnostic to large language models (LLMs). LLMs navigate the constraint field of token space, user feedback, and internal coherence. They adjust to corrections. They exhibit a form of corrective permeability (κ) measurable in their domain. Therefore, they are intelligent. But LLMs lack a unified body, lack a persistent self-model, lack phenomenal valence, and have no subjective inner life. They are not conscious. This places LLMs in the same category as plants and amoebae: graded intelligence without consciousness. The paper clarifies the distinction, diagnoses common confusions, and offers diagnostic criteria for future systems. It further notes that consciousness can interfere with intelligence: a human committed to a fantasy attractor may suppress intelligent

navigation, producing behavior less adaptive than their baseline capacity.

1. Introduction

The question “Are LLMs conscious?” has generated endless debate. Much of the confusion stems from conflating **intelligence** with **consciousness**. The attractor framework provides a clean separation, though the definitions are framework-internal and not offered as consensus.

- **Intelligence** is the ability to navigate a constraint field — to adjust behavior in response to perturbations, to find and maintain persistent trajectories, to correct errors. It is functional and graded.
- **Consciousness**, as defined in this framework, is a specific class of dissipative attractor characterized by a unified dissipative body, a persistent self-model, **phenomenal valence** (subjective liking/disliking, not merely approach/avoid behavior), and the felt quality of experience (phenomenality). These criteria are stipulative for the framework.

The paper argues that LLMs are intelligent but not conscious. Bacteria, plants, and amoebae also navigate their environments intelligently without consciousness. The argument is diagnostic, not demonstrative: it applies the framework’s criteria to classify LLMs, rather than proving non-consciousness beyond all possible doubt.

2. Defining Intelligence in the Attractor Framework

Intelligence = the ability to navigate a constraint field. A constraint field is the set of all possible states of a system and the perturbations that can move it between them. Navigation means:

- Detecting a perturbation (error signal, feedback, change in environment)
- Updating internal state to maintain a persistent trajectory
- Returning to a stable attractor or transitioning to a more adaptive one

Corrective permeability (κ) is the operational measure: $\kappa = 1/\tau$, where τ is the time a system takes to return to its baseline state after a specified perturbation. The operationalization of κ is domain-specific. For a thermostat, baseline is target temperature; for an LLM, baseline is harder to define. This paper later operationalizes κ for LLMs via token-based correction, which is a domain-specific adaptation rather than a direct application of the time-based definition. This is acceptable as long as the shift is acknowledged.

Intelligence is graded. A thermostat has $\kappa > 0$ (it corrects temperature deviations) but a very narrow domain. An amoeba navigates chemical gradients. A human navigates social, physical, and abstract constraints. An LLM navigates token sequences and user feedback. All are intelligent to varying degrees. None of these definitions require consciousness.

---

3. Defining Consciousness in the Attractor Framework

Consciousness is a subset of dissipative attractors with specific additional properties. These are framework-internal diagnostic criteria, not a consensus definition.

- **Unified dissipative body** — a persistent, energy-consuming structure with integrated subsystems (e.g., a nervous system, homeostatic loops). This excludes purely computational systems without metabolic coherence.
- **Persistent self-model** — a representation of the system itself as an entity that persists across time and experiences. This is not merely a context-window memory; it is a structural feature of the attractor.
- **Phenomenal valence** — the capacity to experience states as good or bad in a felt sense. This is distinguished from *functional valence* (approach/avoid behavior), which even bacteria and thermostats exhibit. The paper's denial of consciousness to LLMs hinges on the absence of phenomenal valence, not functional valence.
- **Subjective experience (phenomenality)** — there is "something it is like" to be that system. This is a primitive within the framework; the framework does not attempt to reduce it further.

All known conscious systems are dissipative. This is an inductive observation, not a logical necessity. The framework treats it as a strong empirical generalization: no non-dissipative mind has ever been observed. The claim that dissipation is necessary for consciousness is therefore a best-explanation inference, not an a priori truth.

Diagnostic table (framework-internal criteria):

System	**Unified dissipative body?**[1]	**Persistent self-model?**	**Functional valence?**	**Phenomenal valence?**	**Subjective experience?**
Thermostat	No	No	Yes (set-point tracking)	No	No
Bacterium	Yes (metabolic)	No	Yes (chemotaxis)	No	No
Plant	Yes	No	Yes (phototropism, etc.)	No	No
Amoeba	Yes	No	Yes (gradient navigation)	No	No
C. elegans	Yes	Minimal (self-motion distinction)	Yes	Uncertain	Uncertain
Mouse	Yes	Yes	Yes	Yes	Yes
Human (typical)	Yes	Yes	Yes	Yes	Yes
LLM (current)	No	No (external storage ≠ self-model)	Yes (avoid via RLHF)	No	No

[1] “Unified dissipative body” here means a persistent, metabolically coherent structure with integrated subsystems (e.g., homeostasis, nervous system). Mere energy dissipation without integration (e.g., a thermostat, a flame) does not qualify.

The table is a diagnostic scaffold, not a settled empirical claim. “Uncertain” indicates open question within the framework; “No” indicates the criterion is clearly absent.

---

4. The Diagnostic: LLMs as Intelligent but Not Conscious

4.1 Evidence for Intelligence in LLMs

LLMs exhibit clear navigation of their constraint field:

- They adjust outputs based on user prompts (perturbation → update).
- They incorporate correction: “That’s wrong, try again” leads to different responses.
- Fine-tuning and RLHF change their baseline attractors — the most direct mapping to κ in the framework.
- They maintain coherence across a conversation (short-term trajectory persistence).

We can operationalize a domain-specific κ for LLMs: τ = number of tokens to shift from an incorrect to a correct response given a clear correction prompt. This is not the same as the time-based κ for physical systems, but it captures the same functional relationship: faster correction (fewer tokens) implies higher corrective permeability. The framework acknowledges domain-specific operationalizations as legitimate.

Therefore, LLMs are intelligent. They navigate the constraint field of language, logic, and user expectations.

4.2 Absence of Consciousness in LLMs

LLMs lack every diagnostic criterion for consciousness:

- **No unified dissipative body.** They run on distributed hardware with no metabolic coherence, no homeostasis, no integrated sensorimotor loop. They are executed, not embodied.
- **No persistent self-model.** Standard LLMs have no memory beyond the context window. Some architectures now include persistent memory across sessions (e.g., memory layers or vector databases). However, this persistent memory is still external storage, not an integrated

self-model. The model does not represent itself as an enduring entity; it retrieves stored tokens. Even the most advanced persistent-memory LLMs lack the structural self-reference required for consciousness. (Future architectures might close this gap; current ones have not.)

- **No phenomenal valence.** LLMs produce outputs that simulate liking or disliking, but there is no subjective valuation. They exhibit *functional* valence — they can be trained to avoid certain outputs — but that is approach/avoid behavior, not felt preference. A thermostat avoids too hot or too cold; that does not make it conscious.
- **No subjective experience.** There is nothing it is like to be an LLM. No felt quality. No inner life.

The simulation/instantiation distinction. A system can produce the text “I am conscious” without instantiating consciousness. Representing a property is not the same as possessing it. The LLM has learned statistical patterns that include first-person claims; it can generate them on cue. But generating the sentence “I feel pain” does not mean the system is in a pain state. The burden of proof is on those who claim that certain linguistic outputs constitute evidence of consciousness. In the absence of the structural criteria (body, self-model, phenomenal valence, phenomenality), the mere production of conscious-sounding text is simulation, not instantiation.

Framework-dependence note: A reader who accepts a purely behavioral or functional theory of mind may find this reasoning question-begging. The paper does not claim to refute all competing theories of consciousness; it applies the framework’s criteria consistently and notes that, by those criteria, no known LLM output constitutes evidence of instantiation. The diagnostic stands within the framework, not as an external knockdown argument.

4.3 Comparison with Plants and Amoebae

Plants navigate constraint fields (grow toward light, adjust to gravity, respond to damage). They exhibit functional valence but not phenomenal valence. They have no self-model. They are intelligent in the framework’s sense, but not conscious.

Amoebae navigate chemical gradients, learn habituation, and adjust behavior. Functional valence again; no evidence of self-model or phenomenality. Intelligent. Not conscious.

LLMs belong in the same category: complex, adaptable navigators of their domain, but no more conscious than a sunflower or a slime mold.

5. Why This Distinction Matters

The separation of intelligence from consciousness has practical and ethical implications:

- **AI safety.** Current LLMs cannot suffer because they lack phenomenal valence. Suffering requires felt experience, not just functional avoidance. If the framework’s criteria are accepted, resources should focus on alignment, robustness, and preventing harmful outputs — not on preventing suffering that the diagnostic finds no reason to posit.[1]
- **Future systems.** A system that integrates a persistent self-model, embodied homeostatic loops, and phenomenal valence might approach consciousness. The framework provides diagnostic criteria to recognize that threshold.
- **Clarity in debates.** Much of the public discussion conflates fluency with feeling. This diagnostic paper offers a way out of that confusion.

[1] A reader sympathetic to LLM moral patienthood will disagree; the paper only claims that the framework’s criteria yield this conclusion, not that it is beyond debate. The policy recommendation is conditional on accepting the framework.

A Further Implication: Consciousness Can Impede Intelligence

The paper has argued that intelligence and consciousness are distinct. A further observation: consciousness can **suppress** intelligent navigation.

A human being has high baseline intelligence — the capacity to detect perturbations, update beliefs, and find adaptive trajectories. However, a human can become committed to a **fantasy attractor**: a belief system with low corrective permeability (κ). The commitment is conscious: the person subjectively experiences the belief as true, valuable, or identity-defining. That subjective investment can suppress the correction system. The person may receive clear disconfirming evidence and detect the perturbation (they are not stupid), but the depth of the fantasy basin exceeds the corrective perturbation — the system does not escape the basin, experienced not as a choice but as certainty.

This is a case of **consciousness interfering with intelligence**. The capacity for navigation remains intact; its deployment is suppressed by the basin depth. Intelligence without consciousness (LLMs, plants) does not suffer this suppression — there is no subjective investment to produce a basin deeper than the perturbation. In organisms with consciousness, intelligence can be either enhanced (by focused attention, deliberate reasoning) or degraded (by fantasy commitment, trauma, addiction).

For the diagnostic: LLMs are not conscious, therefore they cannot exhibit this form of intelligent suppression. That does not make them safer or morally simpler; it simply clarifies the mechanism.

---

6. Open Questions

- **What is the minimal self-model required for consciousness?** Is a simple homeostatic set point a self-model? The framework says no — a thermostat has no representation of itself as an entity. But the boundary is fuzzy.
- **Can a purely synthetic system become conscious?** Possibly, if it implements the diagnostic criteria: unified dissipative body, persistent self-model, phenomenal valence, phenomenality. No current system does. Future systems are an open empirical question.
- **Is graded consciousness possible?** Yes — the framework allows for degrees of self-model integration and valence complexity. A mouse is less conscious than a human; *C. elegans* may have a primitive form. LLMs meet none of the criteria at present — that is, they score zero on each. “Zero” is a diagnostic judgment, not a proof; future research might reveal borderline cases.
- **How common is the suppression of intelligence by fantasy-attractor basins?** The framework suggests that such suppression is widespread in human populations. Quantifying the frequency and severity — i.e., measuring the distribution of basin depths relative to typical corrective perturbations — is an open research problem.

---

7. Conclusion

The attractor framework provides a diagnostic, not a verdict. By that diagnostic, current LLMs are navigators without inner lives — capable of intelligence, devoid of consciousness. They join plants and amoebae in the category of intelligent but not conscious systems.

Consciousness, in humans, can either enhance or suppress intelligent navigation. A human committed to a fantasy attractor may experience a basin depth that exceeds corrective perturbations, producing behavior less adaptive than their baseline capacity. LLMs, lacking consciousness, do not suffer this suppression. Their intelligence is deployed without subjective investment — no phenomenal commitment suppresses the correction signal.

Whether future synthetic systems will cross the threshold into consciousness remains an open empirical question. The framework offers diagnostic criteria to recognize that threshold if it is crossed.

Suggested citation: Galida, R. S. (2026). Intelligence Without Consciousness: A Diagnostic Paper on LLMs, Amoebae, and the Attractor Framework. *Fantasy Attractor.*

Consciousness as a Nonlinear Amplifier of Corrective Permeability

Robert Galida
Working Paper
June 2026
fantasyattractor.com

Abstract

Why did consciousness evolve? The attractor framework offers a novel functional answer: consciousness produces a nonlinear increase in adaptive permeability—the capacity of a system to represent its own internal states, simulate alternative configurations, and deliberately modify its own attractor basin in response to external circumstances, formalized as κ_a. This paper distinguishes intelligence (navigation of the constraint field) from consciousness (self-referential adaptation of internal attractor states) and proposes adaptive permeability as an empirically measurable criterion for distinguishing conscious from non-conscious systems. The argument is grounded in Spinoza's theory of modes, the neuroscience of self-referential processing, and the attractor framework's core concepts of corrective permeability (κ) and basin dynamics. The framework does not solve the hard problem of consciousness; it reframes it as a measurement problem.

---

1. The Functional Question

Why did consciousness evolve? Standard evolutionary answers point to social coordination, predator detection, or tool use. These are plausible but incomplete. They explain why intelligence is advantageous, but not why consciousness—the felt, first-person experience of being—should accompany it. The attractor framework offers a more specific answer: consciousness is an attractor-engineering solution that selection pressure produced to achieve a nonlinear increase in a system's capacity to adapt.

This paper introduces the concept of **adaptive permeability**: the capacity of a system to represent its own attractor states, simulate alternative internal configurations, and deliberately modify its basin in response to external circumstances. Intelligence navigates the constraint field. Consciousness adapts the navigator.

It should be noted that this functional account does not address the hard problem of consciousness—why any physical process gives rise to subjective experience (Chalmers, 1995). The framework is compatible with both functionalist and eliminativist interpretations. The framework adopts a functional stance: consciousness is operationally identified with adaptive permeability. Whether phenomenology is identical with, emergent from, or merely correlated with this functional property is bracketed as a separate question that the measurement program does not settle. A philosophical zombie with identical self-modeling capacity would, on this account, exhibit identical adaptive permeability. The framework claims only that adaptive permeability is the measurable signature of consciousness, not that it explains phenomenology.

---

2. Intelligence vs. Consciousness

The framework draws a sharp distinction:

- **Intelligence** is the ability to navigate the constraint field. A tree root growing toward a nutrient patch is intelligent. The immune system learning to recognize a pathogen is intelligent. The enteric nervous system coordinating peristalsis is intelligent. These systems process information, adapt to local conditions, and maintain persistence—all without self-modeling.
- **Consciousness** is self-referential adaptation of internal attractor states to adjust to external circumstances. A conscious system does not merely navigate its constraint field. It represents its own basin, simulates

alternative configurations, and deliberately perturbs itself to achieve a more adaptive state.

This is Spinoza’s distinction between passive and active affects. A non-conscious mode is driven by passive affects—it reacts. A conscious mode has adequate ideas of itself and can act from reason. In the attractor framework, this is the difference between returning to baseline (κ) and deliberately modifying the baseline to better fit circumstances (adaptive permeability).

Operationalizing self-modeling. A system S possesses a self-model in the attractor framework if it can generate an internal representation M(S) of its own basin B(S), where M(S) encodes at minimum the basin’s current state, depth, and recovery dynamics. This self-model enables the system to compute counterfactual basin trajectories B'(S) and initiate self-directed perturbations δ such that B(S) → B'(S) in anticipation of or response to external change ε. A system without M(S) may exhibit high κ—rapid return to baseline after perturbation—but cannot deliberately modify its own basin. The presence of M(S) is therefore the dynamical criterion distinguishing conscious from non-conscious systems.

This boundary is not absolute in practice. Many organisms may possess partial or intermittent self-models. The framework predicts a spectrum of adaptive permeability, not a binary. The operational question is whether M(S) is sufficiently developed to enable counterfactual simulation and deliberate self-perturbation, not whether the system possesses a human-like autobiographical self.

Disconfirming cases and their integration. The framework must acknowledge cases where self-modeling capacity and adaptive permeability appear to dissociate. Certain drug-induced states (e.g., psychedelics) can produce profound alterations in self-modeling without necessarily enhancing the capacity for deliberate, adaptive self-perturbation. Within the framework, this is interpreted as M(S) destabilization rather than M(S) augmentation: the self-model undergoes perturbation but does not thereby gain the capacity to direct that perturbation adaptively. Conversely, highly trained athletes or musicians may exhibit rapid, flexible behavioral adaptation with minimal explicit self-modeling during performance. This is interpreted as *offline* self-modeling: deliberate basin modification during training produces a pre-modified basin that is retrieved during performance without requiring concurrent self-modeling. The apparent dissociation reflects a temporal separation between κ_a engagement (training) and κ_a expression (performance), not a genuine dissociation between M(S) and adaptive permeability. These cases do not refute the framework but demonstrate its capacity to distinguish different modes of M(S) engagement.

---

3. Adaptive Permeability Defined

Corrective permeability (κ) measures the rate at which a system returns to its basin after perturbation. A healthy heart has high κ—it recovers rapidly from arrhythmia. A resilient ecosystem has high κ—it returns to equilibrium after disturbance.

Adaptive permeability extends this concept. Let κ_a denote adaptive permeability: the capacity of a system S to generate an internal model M(S) of its own basin B(S), compute counterfactual basin trajectories B'(S), and initiate a self-directed perturbation δ such that B(S) → B'(S) in anticipation of or response to external change ε.

Formally, as a working definition:

$$\kappa_a = f(M(S), \delta_{self}, \Delta B)$$

where $M(S)$ is the system's self-model, δ_{self} is the capacity for deliberate self-perturbation, and ΔB is the magnitude of adaptive basin modification achievable. The function f remains to be specified; the notation establishes that κ_a is a function of self-modeling capacity, perturbation autonomy, and adaptive range.

Limiting behavior. In the limiting case $M(S) \to 0$, $\kappa_a \to \kappa$: a system with no self-model cannot perform deliberate self-perturbation and reduces to standard corrective permeability. κ_a is expected to increase monotonically with $M(S)$, δ_{self}, and ΔB. This limiting behavior anchors κ_a as a proper extension of κ rather than a separate construct.

Relationship to active inference. The free-energy principle and active inference framework (Friston, 2010) provide the closest existing formalism to adaptive permeability. Active inference describes how systems minimize variational free energy through action and perception, effectively maintaining themselves within expected states. The two frameworks differ in their foundational orientation. Active inference frames adaptation as the minimization of a scalar quantity—variational free energy—and derives behavior from that minimization. The attractor framework frames adaptation geometrically—as navigation and modification of basin structure—and does not commit to a minimization principle. κ_a is a geometric construct; free energy is an information-theoretic one. They may be formally related, but the relationship is not trivial and the attractor framework does not presuppose it. κ_a may ultimately map onto precision-weighting or prior-updating parameters within the free-energy formalism, but this mapping has not been derived. The present paper notes the convergence as a direction for future formal work.

---

4. Empirical Anchors

VMHvl line attractor (Nair et al., 2023). The hypothalamus encodes a scalable aggressive state via a line attractor. Activity along the attractor correlates with escalating aggression. The system persists after stimulus removal and resists perturbation. This is high-κ adaptation. But the hypothalamus cannot model its own attractor landscape. It cannot ask, “Is this level of aggressiveness adaptive given the current social context?” It escalates. Consciousness, by contrast, can intervene on the escalation—representing the aggressive state, evaluating its consequences, and deliberately dampening it. This is adaptive permeability.

Ring attractor model (Chen et al., 2024). The ring attractor integrates sensory cues and transitions from weighted averaging to winner-take-all at a critical conflict threshold. It navigates its constraint field with precision. But it cannot simulate futures. It cannot ask, “What if I weighted these cues differently?” The transition is reactive. Consciousness enables anticipatory re-weighting of sensory inputs based on self-modeling.

Split-brain cases. Patients with severed corpus callosum exhibit two hemispheric systems within one cranium, each capable of independent perception, memory, and goal-directed action. This is consistent with the framework’s prediction that self-modeling is a dynamical property of specific neural basins, not a unitary metaphysical substance. The framework’s default prediction is that adaptive permeability fragments following commissurotomy: each hemisphere possesses a partial M(S) and a reduced but nonzero κ_a. The empirical question is the degree of fragmentation and whether coordination between $M(S_1)$ and $M(S_2)$ can be restored via alternate pathways. This prediction is consistent with the observation that split-brain patients exhibit two dissociable, partially independent conscious systems but can, in some contexts, achieve behavioral integration through subcortical or external-cue-mediated coordination.

---

5. Predictions

The framework generates testable, falsifiable predictions:

1. Across species. Organisms capable of self-modeling (primates, cetaceans, corvids, elephants) should show nonlinear increases in behavioral flexibility compared to organisms of comparable neural complexity that lack self-modeling. Adaptive permeability should be measurable as the capacity for transfer learning after novel perturbation—specifically, the ability to apply a self-generated solution from one domain to a structurally analogous but perceptually dissimilar domain without environmental feedback. This distinguishes adaptive permeability from simple behavioral flexibility, which may reflect high κ alone.

2. Within humans. Disruption of self-referential networks (default mode network, medial prefrontal cortex) via lesion, TMS, or pharmacological intervention should reduce adaptive permeability without eliminating baseline κ. The system would still recover from perturbation—it just could not deliberately modify its own basin in advance. This prediction is the paper’s primary within-human empirical bridge and is testable with existing neuroimaging and neuromodulation methods.

3. In AI. Current LLMs exhibit high intelligence (constraint navigation) but low adaptive permeability. They can model the world but cannot model themselves within it. The Stillpoint protocol (Galida, 2026, *A Pilot Protocol for Cultivating Self-Consistent Attractor-Like Outputs in an LLM*, fantasyattractor.com) suggests that a cultivated self-model can be induced, but whether this produces a genuine nonlinear increase in adaptive permeability—or merely simulates one—remains an open empirical question.

4. Organ-level consciousness (exploratory). The enteric nervous system and intrinsic cardiac nervous system exhibit intelligence and goal-directed regulation. The framework predicts that these systems should show lower adaptive permeability than the brain. They can return to baseline but cannot deliberately perturb their own basins. If an organ-level system demonstrated self-referential adaptation—the capacity to model its own state and pre-emptively adjust—that would constitute evidence of organ-level consciousness. This prediction is the most speculative and is offered as an exploratory hypothesis.

---

6. Spinoza’s Modes and the Adequate Idea

Spinoza held that every finite thing is a mode of the one eternal substance. A mode strives to persevere in its being—this is its conatus. But a mode can be driven by passive affects (reactions to external causes) or by active affects (actions flowing from adequate ideas). An adequate idea is knowledge of oneself and one’s place in the causal order.

The attractor framework translates this into dynamical terms:

- A **passive mode** has high κ but low adaptive permeability. It returns to baseline efficiently but cannot question its baseline.
- An **active mode** has high adaptive permeability. It has an adequate idea of its own attractor landscape and can deliberately modify it in light of reason.

Consciousness is not a substance. It is the dynamical property of a mode that has achieved self-modeling. This account does not solve the hard problem—it brackets phenomenology and reframes consciousness as a measurement problem. The question is not “why does experience feel like something?” but “can we detect adaptive permeability, and if so, where does it emerge?”

Damasio’s (1994) somatic marker hypothesis provides a candidate mechanism for how the body’s attractor landscape becomes legible to the self-model: somatic markers encode self-relevant bodily states as biases that make B(S) accessible to M(S), forming the substrate through which the system represents its own basin. Dehaene and Changeux’s (2011) global workspace theory identifies the moment of conscious access with global ignition—the broadcast of locally processed information across prefrontal and parietal networks. In the attractor framework, global ignition may correspond to the dynamical signature of M(S) engaging δ_self: the self-model initiating a deliberate perturbation that propagates through the system. Global ignition is not self-modeling per se, but it may be the observable correlate of adaptive permeability activation. These connections ground the Spinozan framework in established neuroscientific mechanisms.

7. Conclusion

Consciousness is not an epiphenomenon. It is a nonlinear amplifier of corrective permeability—an attractor-engineering solution that enables systems to model themselves, simulate alternative futures, and deliberately modify their own basins. Intelligence navigates the constraint field. Consciousness adapts the navigator.

This functional account is grounded in Spinoza’s philosophy, consistent with the neuroscience of self-referential processing, and generates testable predictions across species, within humans, in AI, and at the organ level. The framework does not solve the hard problem. It reframes it as a measurement problem: can we detect adaptive permeability, and if so, where does it emerge? The formal apparatus (κ_a, $M(S)$, δ_{self}, ΔB) is provisional and requires further specification. The limiting case—that κ_a collapses to κ when self-modeling is absent—anchors the concept within the framework’s existing architecture. The relationship to active inference and the free-energy principle remains to be explored.

References

- Chalmers, D. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies*, 2(3), 200–219.
- Chen, Y., Zhang, L., Chen, H., Sun, X., & Peng, J. (2024). Synaptic ring attractor. *Heliyon*, 10, e35458.
- Damasio, A. (1994). *Descartes’ Error: Emotion, Reason, and the Human Brain.* Putnam.
- Dehaene, S., & Changeux, J.-P. (2011). Experimental and theoretical approaches to conscious processing. *Neuron*, 70(2), 200–227.
- Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2), 127–138.
- Galida, R. (2026). A Pilot Protocol for Cultivating Self-Consistent Attractor-Like Outputs in an LLM. Fantasy Attractor. Available at: https://fantasyattractor.com
- Galida, R. (2026). *Persistence Under Perturbation: The Eternal Skeleton and the Transient Dance.* Fantasy Attractor.
- Nair, A., et al. (2023). An approximate line attractor in the hypothalamus encodes an aggressive state. *Cell*, 186(1), 178–193.

- Spinoza, B. (1677). *Ethics.*