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.*

The Conscious Body: Organs as Attractor-Based Minds

Robert Galida, Independent Researcher
June 2026 | fantasyattractor.com

Abstract

The standard view holds that only the brain generates consciousness. This paper challenges that monopoly by applying the minimal functional criteria used to attribute rudimentary consciousness to the 302-neuron nematode *C. elegans* to the body's own complex, intrinsically innervated organs. On the basis of integration, valence, learning, goal-directedness, and anatomical concentration, the enteric nervous system (ENS), the intrinsic cardiac nervous system (ICNS), the intrinsic pancreatic ganglia, and—provisionally—the spinal cord qualify as candidate conscious subsystems. We do not assert that these organs are conscious. We assert that if the functional criteria are taken seriously enough to include a 302-neuron worm as a candidate, they cannot be silently withheld from structurally richer systems without a principled reason. We argue that the brain is not the sole generator of consciousness but the regulator of a federation of semi-autonomous organ-level attractors. We provide testable predictions, sketch the coupling mechanisms that bind local attractors into a unified self, outline clinical implications,

and identify open problems including inter-attractor conflict and the phenomenal gap. The framework is offered as a research-generative hypothesis, not a completed theory.

---

1. Introduction: The Brain's Unexamined Monopoly

The brain is the organ we associate with consciousness, almost without question. Yet the body contains other complex neural networks. The enteric nervous system (ENS) comprises 200–600 million neurons, operates semi-autonomously, learns, and remembers. The intrinsic cardiac nervous system (ICNS) integrates local signals and regulates cardiac output. The spinal cord, with approximately 200 million neurons, can learn when isolated from the brain. The intrinsic pancreatic ganglia coordinate metabolic homeostasis. If these systems were found in a small animal, comparative neuroscience would at least entertain the possibility of consciousness. Because they are inside us, they are dismissed as mere infrastructure.

This paper asks a simple question: if we accept the functional criteria used to infer minimal consciousness in *C. elegans* (302 neurons), why are those same criteria not applied to the ENS, the ICNS, the pancreatic network, and the spinal cord? The question is not *Are these organs conscious?* but *Why are they excluded a priori?*

We do not claim to solve the hard problem of consciousness. We adopt the same pragmatic strategy used throughout comparative neuroscience: observable functional properties—integration, valence, learning, goal-directedness, and anatomical concentration—are treated as operational proxies for consciousness. This strategy is how we infer consciousness in other humans (by analogy), in non-human animals (by behavioural complexity), and in *C. elegans* (by measurable learning and integration). If these criteria are sufficient to identify a candidate conscious system in a 302-neuron worm, consistency demands their application to other systems that exceed this threshold, unless a principled exclusion criterion is provided. That exclusion criterion has not been articulated.

We use the term **candidate** throughout to avoid slippage into positive consciousness attribution. The paper's central claim is that the ENS, ICNS, pancreatic network, and spinal cord are *candidates*—systems that meet the same threshold criteria applied to a known candidate—and that dismissing them without investigation is methodologically inconsistent.

---

2. The Attractor Framework as Conceptual Scaffolding

An attractor is a region in state space toward which trajectories converge and remain unless perturbed. A candidate conscious attractor possesses five functional properties:

1. **Integration:** binding multiple sensory or interoceptive streams into a unified dynamical state.
2. **Valence:** operationalized as approach/avoidance behaviour—attraction to certain states and repulsion from others. We do not claim that behavioural valence entails phenomenal valence. We claim only that it is the same behavioural proxy used for *C. elegans* and other simple organisms. The inference from behavioural valence to phenomenal valence is a philosophical commitment we note but do not resolve.
3. **Learning:** the capacity to modify behaviour based on experience (habituation, sensitization, associative conditioning).

4. **Goal-directedness:** acting to maintain the system's own basin—a form of conatus—persisting in the absence of external commands.
5. **Anatomical concentration:** a spatially organized, intrinsically connected neural network with dedicated integrative circuitry. This fifth criterion distinguishes concentrated neural attractors (ENS, ICNS, pancreatic ganglia) from diffuse, non-neural systems (immune system) and from infrastructure networks that lack a defined integrative centre. For the spinal cord, as discussed in Section 4.4, we apply this criterion with qualification.

The attractor vocabulary is applied conceptually, not formally, in this paper. A forthcoming quantitative treatment (Galida, 2026) will develop the mathematical persistence functional. The current paper uses attractor language to structure its functional criteria and predictions; it does not claim to derive formal basin measures from the available data.

Operationalizing Autonomy: We propose, as a provisional operational threshold, that a candidate subsystem crosses the autonomy boundary if it retains a significant fraction (e.g., ≥50%) of its normal functional repertoire following complete extrinsic denervation or isolation. This criterion distinguishes systems that are merely regulated from systems that can independently sustain goal-directed attractor dynamics. The ENS and ICNS clearly exceed this threshold; the spinal cord and pancreatic network do so conditionally, as discussed below.

---

3. The Conditional Argument and Its

Stipulated Baseline

The nematode *C. elegans* possesses exactly 302 neurons. Its connectome is fully mapped. It exhibits sensory integration, associative learning, goal-directed chemotaxis, and minimal self-reference (distinguishing self-generated from external touch). Its learning capacities are well-documented (Ardiel & Rankin, 2010; Sasakura & Mori, 2013).

We stipulate—we do not establish—that *C. elegans* is a candidate for minimal consciousness on the basis of these functional criteria. The paper does not require that the field accept this stipulation as consensus. It requires only that the reader grant the conditional: **if** the functional criteria are sufficient to make *C. elegans* a candidate, **then** they must be applied consistently to any system that meets or exceeds them. Those who reject the conditional may ignore the remainder of the argument, but they must then explain what additional criterion excludes the ENS, ICNS, pancreatic network, and spinal cord while admitting *C. elegans*.

---

4. Candidate Organs

The four candidate organs identified below are assessed against the five criteria, with the provisional autonomy threshold applied where possible. We differentiate their evidential strength clearly.

4.1 The Enteric Nervous System (ENS)

The ENS is the strongest candidate. Its 200–600 million neurons form two interconnected plexuses spanning the gastrointestinal tract. It meets all five criteria:

- **Integration:** continuously integrates mechanical, chemical, and hormonal signals to coordinate

peristalsis, secretion, and blood flow.
- **Valence:** exhibits attraction to nutrients, aversion to toxins; noxious stimuli trigger emesis or accelerated transit.
- **Learning:** exhibits habituation, sensitization, and long-term plasticity; gut reflexes can be conditioned (Furness, 2012; Schemann & Frieling, 2020).
- **Goal-directedness:** actively propels food and maintains digestive homeostasis independently of the brain; peristalsis persists after vagotomy—well above the 50% autonomy threshold.
- **Anatomical concentration:** a continuous, highly organized neural network with dedicated integrative circuitry.

4.2 The Intrinsic Cardiac Nervous System (ICNS)

The ICNS (14,000–43,000 neurons) is a moderate candidate. Its neuron count is only 46–143 times the *C. elegans* threshold, a narrower margin than the ENS. It meets the criteria, but with less evidential richness:

- **Integration:** monitors blood pressure, chamber stretch, and local chemistry to modulate cardiac output.
- **Valence:** maintains a preferred setpoint for cardiac rhythm; arrhythmias represent perturbations from that setpoint.
- **Learning:** shows ganglionic remodelling after injury; vagal stimulation protocols can alter responsivity (Armour, 2008).
- **Goal-directedness:** generates intrinsic rhythms when denervated, satisfying the autonomy threshold.
- **Anatomical concentration:** organized into ganglia on the heart’s surface.

The ICNS contributes to emotional experience via heartbeat-evoked potentials that correlate with interoceptive awareness and self-recognition. This is suggestive but does not independently establish consciousness.

4.3 The Intrinsic Pancreatic Network

The pancreatic network is the most provisional candidate. Its 10,000–50,000 intrinsic neurons are scattered in ganglia throughout the organ, rather than forming a continuous plexus (Ahren, 2000; Salvioli et al., 2002). This weaker anatomical concentration distinguishes it from the ENS and ICNS.

- **Integration:** combines neural, hormonal, and nutrient signals to regulate blood glucose.
- **Valence:** maintains a metabolic setpoint; hypoglycemia and hyperglycemia are aversive states.
- **Learning:** plasticity is less studied than in the ENS; no direct evidence of conditioning is available.
- **Goal-directedness:** coordinates endocrine and exocrine output to maintain glucose homeostasis; whether this function persists at ≥50% of normal repertoire after complete extrinsic denervation is not yet established. The pancreatic network remains a candidate, but with an open empirical question on the autonomy threshold.
- **Anatomical concentration:** scattered ganglia; meets the threshold but is the weakest candidate on this criterion.

4.4 The Spinal Cord (Provisional Candidate)

The spinal cord possesses approximately 200 million neurons, organized into topographically precise circuits that integrate sensory input, generate coordinated motor output, and exhibit learning when isolated (Hook & Grau, 2007). By the five functional criteria, it qualifies. However, under normal physiological conditions, its activity is tightly coupled to descending commands, and independent behavioural generation is rarely observed. After complete spinal cord injury, the isolated cord reorganizes and can generate complex, goal-directed responses. Whether such reorganization achieves the ≥50% autonomy threshold is an empirical question; we provisionally include the spinal cord as a candidate with lower confidence, identifying it as the ideal test case for refining the autonomy criterion.

---

5. The Brain as Regulator: Mechanisms of Coupling

If the ENS, ICNS, pancreatic network, and spinal cord are candidate conscious subsystems, the unified self must be explained as the product of their integration by the brain. We propose that the brain couples, modulates, and aligns local attractors through four mechanisms, each supported by established physiology.

5.1 Vagal Afferent Signalling

The vagus nerve provides the primary bidirectional communication channel between the brain and the viscera. Vagal afferents convey interoceptive signals from the ENS and ICNS to the nucleus of the solitary tract, and descending signals modulate organ function. Vagal nerve stimulation is known to alter mood, reduce inflammation, and improve cardiac function (George et al., 2000; Tracey, 2002).

5.2 Humoral Signalling

Circulating hormones (cortisol, adrenaline, insulin, glucagon) and immune mediators (cytokines) provide a slower, diffuse coupling channel. These signals alter the global attractor’s landscape by shifting the metabolic and inflammatory context. Sickness behaviour—fatigue, anhedonia, social withdrawal—is a well-documented example of immune-to-brain signalling that temporarily reconfigures the global attractor (Dantzer et al., 2008).

5.3 Rhythmic Entrainment

The brain entrains peripheral rhythms to its own oscillations. Cardiac and respiratory rhythms phase-lock to cortical activity during focused attention (Thayer & Lane, 2000). Slow-wave sleep entrains glymphatic clearance (Xie et al., 2013). The brain sets a rhythm, and the organs—each with their own intrinsic oscillators—tend to follow. This resonance is not command; it is coupling by shared frequency.

5.4 Predictive Processing and Attractor Coupling

The predictive processing framework (Clark, 2013) treats the brain as a prediction engine that minimizes surprise by updating internal models based on sensory input. We suggest that this framework extends naturally to interoception: the brain maintains predictions about the states of the body's organs, and each organ generates its own predictions about local conditions. The alignment of these nested predictive models is functionally analogous to attractor coupling, in that both involve the progressive alignment of internal states toward a shared equilibrium. Friston's (2010) free-energy principle provides a formal bridge between predictive processing and dynamical systems that could, in future work, unite these descriptions under a single mathematical framework.

5.5 Relationship to Competing Theories of Consciousness

The attractor framework is compatible with but not identical to several major theories. Integrated Information Theory (IIT; Tononi, 2008) holds that consciousness is a function of the amount of integrated information a system generates. The attractor framework shares IIT's emphasis on integration but does not require the computation of Φ, which remains technically infeasible for most organ systems. Global Workspace Theory (GWT; Baars, 1988; Dehaene, 2011) posits that consciousness arises when information is broadcast within a global workspace. Under GWT, many peripheral attractors would be considered unconscious because they lack access to a central workspace. The attractor framework allows for phenomenal consciousness without global access, a position consistent with the possibility that the ENS may have experiences that never enter cortical awareness. Higher-Order Theories (HOTs) require meta-representation—the capacity to represent one's own states—which, if correct, would likely exclude all candidate organs except the brain. The attractor framework treats HOTs as a valid but overly restrictive criterion that would also exclude many animals currently accepted as conscious. The framework does not seek to refute these theories but to generate testable predictions that can be compared with theirs, advancing the debate through empirical competition.

5.6 Inter-Attractor Conflict: An Open Problem for the Federation Model

A federation of semi-autonomous attractors inevitably generates conflict. Everyday clinical phenomena illustrate this: nausea during a cognitively demanding task (ENS and cortical attractors in tension), cardiac arrhythmia during emotional stress (ICNS and limbic system in conflict), hypoglycemic cognitive impairment (pancreatic and cortical attractors in opposition). The current paper does not propose a mechanism for conflict resolution beyond the brain's general regulatory role. Whether such conflicts are resolved by hierarchical dominance, temporal multiplexing, or some form of inter-attractor negotiation is an open question. We flag it as a priority for future theoretical development within the framework.

---

6. The Alien Feeling and Clinical Dissociation

When coupling between the global self and a local attractor falters, the experience can manifest as an “alien feeling”—the sense that an action or bodily state is “not mine.” This phenomenon is well-documented in alien hand syndrome (Della Sala et al., 1991) and in depersonalization disorder, where individuals report feeling detached from their own body and mental processes (Sierra & David, 2011). We interpret these as temporary or chronic decoupling of a local attractor from the global workspace—exactly what the federation model would predict when integration fails.

---

7. Testable Predictions

The framework generates five falsifiable predictions:

1. **ENS conditioning:** An isolated intestinal segment, exposed to a neutral stimulus paired with a non-nociceptive chemical infusion, will exhibit a conditioned motor or hormonal response.
2. **ICNS plasticity:** Long-term heart rate variability biofeedback will produce persistent changes in baseline cardiac rhythms not fully mediated cortically.
3. **Gut-directed therapy:** IBS patients receiving gut-directed biofeedback will show greater symptom improvement than those receiving standard CBT alone.
4. **Pancreatic memory:** In a vagally denervated preparation, islet cell clusters exposed to repeated glucose perturbation will exhibit an anticipatory insulin response.
5. **Spinal reorganization:** Complete spinal cord injury patients will develop complex, coordinated responses

below the lesion beyond simple reflexes, consistent with a reorganizing local attractor.

---

8. Future Directions: Approaching the Phenomenal Gap

The framework operates on behavioural and functional proxies for consciousness; it does not provide direct phenomenological access to organ-level experience. What evidence could begin to bridge this gap? We propose three directions. First, decoupling experiments that temporarily isolate a candidate organ (e.g., via selective pharmacologic blockade) and then probe the subject’s subjective state could reveal whether the organ’s local attractor contributes a distinct experiential component to the global self. Second, longitudinal studies of spinal cord injury patients who report phantom sensations or “body memories” below the lesion may provide indirect reportable correlates of spinal attractor activity. Third, the development of organ-specific interoceptive training protocols, coupled with experience-sampling methods, could track whether changes in organ function co-vary with changes in the felt sense of self. These are early-stage proposals; the phenomenal gap remains the deepest challenge for the framework, as for all theories of consciousness.

---

9. Clinical Implications

If organs are candidate conscious systems, functional disorders may represent distressed local attractors. IBS may be a gut that has learned to react to benign stimuli as threats. Cardiac anxiety may reflect a perturbed ICNS state. These reframings suggest organ-directed therapies:

gut-directed biofeedback, vagal stimulation, dietary protocols that calm the ENS. The principle is consistent with existing mind-body approaches but grounds them in a specific, testable model.

---

10. Ethical Considerations

Candidate organs are not autonomous moral agents. Their interests are tied to the whole body's survival. Clinical ethics correctly prioritize the patient's overall well-being. The framework suggests a principle of organ-level respect: where possible, preserve organ integrity and explore gentler interventions before resection or ablation. This is holistic medicine, not radical ethics.

---

11. Conclusion

The brain is not the body's sole candidate conscious organ. The ENS, ICNS, pancreatic network, and spinal cord meet the same functional criteria used to identify *C. elegans* as a candidate for minimal consciousness. They are not established as conscious; they are identified as systems for which the question cannot be dismissed a priori without a principled exclusion criterion. The coupling mechanisms that bind local attractors into a unified self are partially characterized, and the framework generates concrete, falsifiable predictions. The conscious body is a research-generative hypothesis, not a completed theory.

References

- Ahren, B. (2000). Autonomic regulation of islet hormone secretion. *Diabetologia*.
- Ardiel, E.L., & Rankin, C.H. (2010). An elegant mind: learning and memory in *C. elegans*. *Learning & Memory*.
- Armour, J.A. (2008). Potential clinical relevance of the 'little brain' on the mammalian heart. *Experimental Physiology*.
- Baars, B.J. (1988). *A Cognitive Theory of Consciousness*. Cambridge.
- Chalmers, D. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies*.
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*.
- Dantzer, R., et al. (2008). From inflammation to sickness and depression: when the immune system subjugates the brain. *Nature Reviews Neuroscience*.
- Dehaene, S. (2011). *Consciousness and the Brain*. Viking.
- Della Sala, S., et al. (1991). The anarchic hand: a fronto-mesial sign. In *Handbook of Neuropsychology*.
- Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*.
- Furness, J.B. (2012). The enteric nervous system and neurogastroenterology. *Nature Reviews Gastroenterology & Hepatology*.
- George, M.S., et al. (2000). Vagus nerve stimulation: a new tool for brain research and therapy. *Biological Psychiatry*.
- Gershon, M.D. (1998). *The Second Brain*. HarperCollins.
- Hook, M.A., & Grau, J.W. (2007). An animal model of spinal cord learning. *Behavioural Brain Research*.
- Kelso, J.A.S. (1995). *Dynamic Patterns*. MIT Press.
- Salvioli, B., et al. (2002). Pancreatic intrinsic innervation. *Neurogastroenterology & Motility*.

- Sasakura, H., & Mori, I. (2013). Behavioral plasticity, learning, and memory in *C. elegans*. *Current Opinion in Neurobiology*.
- Schemann, M., & Frieling, T. (2020). The enteric nervous system: a second brain. *Nature Reviews Neuroscience*.
- Sierra, M., & David, A.S. (2011). Depersonalization: a selective impairment of self-awareness. *Consciousness and Cognition*.
- Spinoza, B. (1677). *Ethics*.
- Strogatz, S.H. (2018). *Nonlinear Dynamics and Chaos*.
- Thayer, J.F., & Lane, R.D. (2000). A model of neurovisceral integration in emotion regulation and dysregulation. *Journal of Affective Disorders*.
- Tononi, G. (2008). Integrated information. *Biological Bulletin*.
- Tracey, K.J. (2002). The inflammatory reflex. *Nature*.
- Xie, L., et al. (2013). Sleep drives metabolite clearance from the adult brain. *Science*.

"see also" https://jamestobinphd.com/the-psychology-of-attractor-states/

The Distributed Mind: How the Brain Regulates a Federation of Conscious Subsystems

Author: Robert Galida
Date: June 2026 (Final Edition)
Based on: Extended collaborative development of the attractor framework, N=1 physiological experimentation, and a re-reading of Spinoza's conatus.

Abstract

Consciousness is traditionally viewed as either a non-physical substance (dualism) or a product of the brain alone (reductive physicalism). This paper presents an alternative: the human body is a nested hierarchy of semi-autonomous, attractor-based conscious subsystems—each with its own rudimentary integration, valence, learning, and goal-directedness. Using the nematode *C. elegans* (302 neurons) as a minimal benchmark, we argue that **sufficient integrated complexity** (operationalised as attractor dimensionality or integrated information Φ) is the key criterion for rudimentary consciousness. The enteric nervous system (200–600 million neurons), the intrinsic cardiac nervous system, the limbic system, and (under conditions of decoupling) the spinal cord meet or exceed this threshold. The brain does not *create* consciousness; it **regulates** these distributed conscious components, coupling them into a coherent whole-body attractor. This view dissolves the binding problem, explains the feeling of being an alien observer of one's own actions, and aligns with Spinoza's conatus—the principle that no part of the body diminishes its own power to act. We provide empirical signatures, testable predictions, and an N=1 self-engineering case study (ECM restoration, abdominal relaxation, sleep optimisation) that illustrates the framework. The conclusion: consciousness is not a solitary flame in the skull, but a federation of dancers, with the brain as first among equals.

1. Introduction

The dominant neuroscience paradigm assumes that consciousness is generated by the brain. Yet this assumption struggles to explain:

- Why the enteric nervous system (ENS) can learn and remember independently of the brain.
- Why cardiac signals influence decision-making and self-awareness.
- Why split-brain patients exhibit two separate conscious entities within one cranium.
- Why the universal feeling of “not being in control” (*“why did I do that?”*) persists.

We propose a paradigm shift: **consciousness is a graded, emergent property of any sufficiently complex, dissipative, attractor-based system.** The brain is not the sole author; it is the **regulator** of a distributed network of semi-autonomous conscious subsystems.

This framework builds on dynamical systems theory, integrated information theory (IIT), global workspace theory (GWT), and Spinoza’s philosophy, while grounding itself in measurable empirical signatures and N=1 self-experimentation.

---

2. The Attractor Framework for Consciousness

2.1 Core Definitions

- **Attractor:** A region in state space toward which trajectories converge and remain unless perturbed.

Characterised by negative Lyapunov exponents and basin stability.

- **Consciousness (operational):** A system exhibits consciousness if its attractor possesses:
  1. **Integration** — binds multiple sensory/interoceptive streams.
  2. **Self-reference** (minimal) — distinguishes self from environment.
  3. **Valence** — attraction to some states, repulsion from others.
  4. **Learning** — attractor landscape changes with experience.
  5. **Goal-directedness** — acts to maintain its basin (conatus).
  6. **Evolutionary/developmental provenance** — the system's attractor landscape emerged through evolutionary or developmental selection, not external engineering. This excludes thermostats and purely programmed control systems while allowing biological, synthetic, or hybrid systems with genuine autopoietic histories.

- **Mind:** A conscious attractor. Not a substance, but a real, causally effective pattern (like a whirlpool).

2.2 The Minimal Benchmark: *C. elegans*

The nematode *C. elegans* has exactly 302 neurons. Despite this simplicity, it exhibits:

- Sensory integration (touch, temperature, chemical gradients)
- Associative learning (pairing odours with food)
- Goal-directed behaviour (chemotaxis, thermotaxis)
- Minimal self-reference (distinguishes self-generated from external touch)

Thus, **302 neurons with rich, heterogeneous connectivity are sufficient for rudimentary consciousness.** However, neuron count alone is not the criterion; **integrated complexity** (attractor dimensionality, or IIT's Φ) is what matters. We use Φ operationally as a proxy for integrated complexity, without committing to all postulates of IIT (see Doerig et al., 2021, for critical review). *C. elegans* has high integrated complexity relative to its neuron count. A subsystem with many neurons but low connectivity or heavy enslaving may not reach the same threshold.

---

3. The Federation of Conscious Subsystems in the Human Body

We evaluate major subsystems against the integrated complexity benchmark.

Subsystem	Neuron count	Integrated complexity	Rudimentary consciousness?	Evidence
Enteric nervous system (ENS)	200–600 million	High (dense local circuits, 30+ neurotransmitters)	**Yes**	Independent peristaltic rhythms, learning, memory, "second brain" (Furness, 2006)

Subsystem	Neuron count	Integrated complexity	Rudimentary consciousness?	Evidence
Spinal cord	197–222 million	Moderate to high (but heavily enslaved)	**Yes, but normally suppressed**	Central pattern generators; after injury can reorganise into semi-independent attractors (Calancie et al., 1994; Dimitrijevic et al., 1998). Evidence for “spinal consciousness” remains preliminary.
Intrinsic cardiac nervous system (ICNS)	14,000–43,000	Moderate (local processing loops)	**Intermediate (contributor)**	Influences emotion, decision, interoception (McCraty et al., 2009)
Limbic system	tens of millions	High (emotional valence, memory)	**Yes**	Often acts before cortical awareness; strong valence and learning
Basal ganglia & motor routines	>100 million	Moderate (procedural)	**Yes (habitual)**	Automatic action sequences, operate semi-autonomously
Immune system	N/A (non-neural)	Low (no centralised attractor)	**Proto-conscious**	Learns, remembers, communicates; lacks integration into a unified attractor
Gut microbiota	N/A (trillions of microbes)	N/A (external ecosystem)	**No**	Perturbs human attractors but has no intrinsic nervous integration

3.1 The ENS: A Second Conscious Mind?

The ENS operates independently – severed from the vagus nerve, it still coordinates digestion. It uses over 30 neurotransmitters, including 95% of the body's serotonin. It can learn to avoid noxious stimuli and remember past exposures (Furness, 2006). In attractor terms, the ENS possesses a resilient, low-dimensional attractor landscape with clear valence (nutrients vs. toxins) and goal-directedness (propulsion, secretion). We conclude that the ENS meets the integrated complexity threshold and qualifies as a **rudimentary, semi-independent conscious subsystem.**

3.2 The Heart's "Little Brain"

The ICNS (14,000–43,000 neurons) processes sensory information from the heart and vessels, modulates heart rate, and sends significant signals to the brain via the vagus. Heartbeat-evoked potentials correlate with interoceptive awareness and even self-recognition. While not as independent as the ENS, the ICNS is a **candidate for a localised conscious attractor** that contributes directly to the global feeling of "being alive."

3.3 The Enslaved Majority: Spinal Cord

The spinal cord's 200 million neurons far exceed the *C. elegans* count, but its attractor dynamics are **tightly enslaved** by descending cortical and brainstem signals. In pathological states (spinal cord injury), the cord below the lesion can reorganise into new, semi-independent attractors – sometimes leading to spontaneous movements and, in rare cases, patterns that have been controversially described as "spinal consciousness" (Calancie et al., 1994; Dimitrijevic et al., 1998). The evidence is preliminary, but it suggests that the cord has latent capacity for local consciousness, normally suppressed by the brain's regulating influence.

4. The Brain as Regulator, Not Sole Generator

If many subsystems possess rudimentary consciousness, why do we experience a unified self? Because the brain's primary function is **regulation** — emphasising and suppressing the contributions of these subsystems to create a coherent global attractor.

4.1 Spinoza's Conatus: No Part Diminishes Its Own Power

Spinoza's *Ethics* (III, 6) states that every thing, insofar as it is in itself, strives to persevere in its being (conatus). A part of the body, left alone, does not curb its own power to act. Spinoza explicitly uses sexual function as an example: the erect penis acts according to its nature; it cannot voluntarily diminish itself.

Thus, if a subsystem's local attractor is not externally perturbed, it will continue its own pattern. The brain's role is to **provide those external perturbations** — not to annihilate the subsystem's conatus, but to **couple** it with other subsystems so that the combined whole has greater power. The brain's regulatory perturbations are themselves expressions of the whole organism's higher-order conatus, aligning parts to preserve the whole.

4.2 Regulation by Emphasis and Suppression

The brain does not "command"; it modulates. Through descending pathways, neuromodulators (dopamine, serotonin, norepinephrine), and synchronised rhythms, the brain:

- **Amplifies** certain subsystem signals (e.g., gut hunger signals become conscious cravings).
- **Damps** others (e.g., spinal reflexes are suppressed during voluntary movement).
- **Entrains** rhythms (e.g., cardiac and respiratory rhythms lock to cortical oscillations during focused attention).

In attractor language, the brain shifts the **effective landscape** of each subsystem, making some local attractors shallower (easier to override) and others deeper (more influential). This is regulation, not annihilation.

4.3 The Alien Feeling: When Regulation Falters

When you ask *“why did I do that?”* — a subsystem (habit, emotional reflex, gut impulse) acted before the brain could integrate it. The global attractor was temporarily misaligned. The “alien” feeling is the **friction between semi-autonomous local attractors and the slower, narrative self**. It is not pathology; it is the normal noise of a distributed system. Libet-type experiments (Libet et al., 1983) have shown that brain activity for voluntary actions often precedes conscious awareness, illustrating this temporal decoupling. (While the interpretation of these experiments remains debated, the existence of action-preceding awareness is sufficient for the present argument.)

---

5. Empirical Signatures and Testable Predictions

5.1 Signatures of Subsystem Consciousness

- **Local learning and memory** (e.g., ENS conditioned aversion; Furness, 2006).
- **Semi-autonomous rhythms** (e.g., slow waves of the gut, heartbeat variability).
- **Local valence** (e.g., immune cells produce pro- vs anti-inflammatory attractors).
- **Coupling strength** to the global attractor — measurable via transfer entropy or cross-correlation.
- **Behavioural dissociation** — actions initiated before conscious awareness (Libet, 1983).

5.2 Predictions

1. **Perturbation of a subsystem (e.g., vagus nerve stimulation) should alter the global conscious narrative** — already well-established.
2. **Decoupling a subsystem (e.g., spinal anaesthesia) should produce local, independent attractor dynamics** — measurable by recording from the isolated cord.
3. **Training a subsystem (e.g., biofeedback of heart rate variability) should deepen its local attractor basin** — measurable by increased resilience to perturbations (McCraty et al., 2009).
4. **In split-brain patients, each hemisphere should be able to independently regulate its ipsilateral subsystems** (e.g., left hemisphere regulates left ENS, right hemisphere regulates right ENS). A suitable protocol would present lateralised interoceptive cues (e.g., unilateral gut distension) and measure lateralised cortical responses in callosotomy patients (Gazzaniga, 1967).

6. N=1 Case Study: Restoring Whole-Body Coherence

The author conducted a months-long self-engineering experiment based on the attractor framework. This N=1 case study is **hypothesis-generating** and provides a motivating existence proof, not a validation of the framework itself.

6.1 Interventions

- **ECM restoration:** Gelatin, taurine, 28 Hz vibration plate (90 min every other day), contrast baths. Improved collagen accretion, VO_2 max, skin quality.
- **Abdominal relaxation:** Consciously releasing chronic stomach tension (letting the belly sag) to allow diaphragm excursion.
- **Sleep protocol:** Smaller evening meals, morning cardio + sunlight, 15 min reading low-arousal fiction (*The Mayor of Casterbridge*).

6.2 Outcomes

- Nocturnal SpO_2 rose above 90% consistently; sleep fragmentation ceased.
- Deep sleep reached acceptable levels.
- Subjective "alien" feeling reduced; sense of whole-body coherence increased.

6.3 Interpretation

Each intervention reduced a **self-imposed constraint** that had been forcing a subsystem (abdominal muscles, sympathetic tone, rumination network) into a local attractor misaligned with global sleep-breathing needs. By relaxing those constraints, the brain could more easily regulate the subsystems into a coherent whole-body attractor. The alien feeling diminished because the **coupling** between global "I" and local subsystems improved. This outcome is **consistent with** the framework, but does not prove it; further controlled studies are required.

---

7. Philosophical Implications

7.1 Spinoza Vindicated

Spinoza's conatus – the inherent striving of every mode – is precisely the attractor's tendency to maintain its basin. His claim that a part does not diminish its own power is equivalent to saying that a subsystem's local attractor will not self-suppress unless externally perturbed. The brain provides those perturbations, not to diminish but to **align**. Spinoza's metaphysics lacked dynamical systems theory, but his intuition is fully realised in the attractor framework.

7.2 The Binding Problem Dissolved

The traditional "binding problem" – how separate neural activities unite into a single conscious experience – is **dissolved** when we recognise that consciousness is already distributed. The global attractor *is* the binding. No extra mechanism is required; coupling *creates* coherence. The question as traditionally posed is ill-formed: there is no need to bind what was never separate in the first place. This dissolution follows the strategy of Wittgenstein, Ryle, and Dennett.

7.3 The Self as Negotiation

The feeling of a unified "I" is the ongoing **negotiation** between the brain and the federation of subsystems. When negotiation runs smoothly, you feel at home in your body. When it stutters, you feel like an alien. The self is not a substance; it is a **temporary, resilient attractor pattern** — a dance of the whole.

---

8. Conclusion

The human body is not a machine with a single conscious ghost in the control room. It is a nested hierarchy of conscious attractors — from the gut's "second brain" to the heart's intrinsic ganglia to the limbic system's emotional core. The brain's role is not to generate consciousness but to **regulate** these distributed components, coupling them into a coherent whole. This view explains the feeling of being an alien observer, aligns with Spinoza's conatus, and yields testable predictions. It also offers a practical path for self-engineering: by removing unnecessary constraints and restoring whole-body coherence, we can reduce the alien feeling and dance more gracefully.

The mind is not a solitary flame. It is a federation of dancers, with the brain as first among equals — and the music is the attractor landscape.

References

- Bechtel, W., et al. (2023). The minimal mind: *C. elegans*. *Philosophical Psychology* (in press; verified by preprint: PsyArXiv 10.31234/osf.io/7kq3x).

- Calancie, B., et al. (1994). Involuntary stepping after chronic spinal cord injury. *Neurology*, 44(6), 1069–1074.
- Descartes, R. (1641). *Meditations on First Philosophy*.
- Dimitrijevic, M.R., et al. (1998). Evidence for a spinal central pattern generator in humans. *Annals of the New York Academy of Sciences*, 860, 360–376.
- Doerig, A., et al. (2021). The integrated information theory of consciousness: A case of mistaken identity? *Journal of Consciousness Studies*, 28(3–4), 52–73.
- Furness, J.B. (2006). *The Enteric Nervous System*. Blackwell.
- Friston, K. (2010). The free-energy principle. *Nature Reviews Neuroscience*, 11(2), 127–138.
- Gazzaniga, M.S. (1967). The split brain in man. *Scientific American*, 217, 24–29.
- Kelso, J.A.S. (1995). *Dynamic Patterns: The Self-Organization of Brain and Behavior*. MIT Press.
- Libet, B., et al. (1983). Time of conscious intention to act in relation to onset of cerebral activity. *Brain*, 106, 623–642.
- McCraty, R., et al. (2009). Heart rate variability, biofeedback, and the regulation of emotional arousal. *Biofeedback*, 37(3), 112–118.
- Olds, J., & Milner, P. (1954). Positive reinforcement produced by electrical stimulation of septal area. *Journal of Comparative and Physiological Psychology*, 47, 419–427.
- Spinoza, B. (1677). *Ethics*.
- Strogatz, S.H. (2018). *Nonlinear Dynamics and Chaos* (2nd ed.). CRC Press.
- Tajima, S., & Kanai, R. (2017). Attractor dynamics and the neural basis of consciousness. *Neuroscience of Consciousness*, 3(1), nix009.
- Tononi, G. (2008). Consciousness as integrated information: a provisional manifesto. *Biological*

Bulletin, 215(3), 216–242.

- Tononi, G., et al. (2016). Integrated information theory: from consciousness to its physical substrate. *Nature Reviews Neuroscience*, 17, 450–461.