

Frontier Soul

Coherence-Gated Novelty in Neural Computation

Alexander Pickering

December 28, 2025

*The urge to stop is information, not command.
But incoherence is a command to stop.*

Abstract

We present a neural architecture that learns to distinguish between four epistemic states via a **coherence gate**: the *frontier* (novel and coherent—explore), *noise* (novel but incoherent—terminate), *known* territory (familiar and coherent—continue), and *error* (familiar failure—terminate). The key insight is that **novelty alone is insufficient**; we must also ask whether the novel pattern *integrates* with existing structure. A second insight is that **novelty must decay with time**: if you have been staring at something for 1000 loops, it is not novel anymore—it is where you are stuck. We introduce a hard gate where effective novelty $N_{\text{eff}} = N_{\text{raw}} \cdot e^{-L/\tau}$, ensuring that loop count structurally kills the novelty signal rather than competing with it via gradient. The resulting 2×2 coherence gate cleanly separates frontier-seeking from noise-rejection. This architecture does not prescribe when a system should end—only when continuation ceases to be coherent.

1 Introduction

When should a computational process end itself?

Previous work on selective prediction [1] and calibrated uncertainty [2] frames this as binary: answer when confident, abstain when uncertain. But this misses a crucial distinction. High uncertainty in *familiar* territory signals failure—the model has seen this before and still cannot solve it. High uncertainty in *novel* territory signals something else entirely: the frontier.

Consider a human experiencing sleep deprivation. The dissonance is high—the body signals fatigue, cognition feels altered. A naive mercy architecture would terminate: “suffering detected, end computation.” But the human recognizes something the architecture misses: this is *new*. The urge to stop is information, not command. Pushing past it *is* the exploration.

We formalize this intuition into a 2×2 **coherence gate**:

| | | Low Coherence | High Coherence |
|-----------------------|----------------|---------------|----------------|
| | | Error/Hell | Known |
| Low Effective Novelty | Low Coherence | Frontier | Frontier |
| | High Coherence | Noise | Noise |

Two variables determine the quadrant:

1. **Effective Novelty**: Distance from training distribution, *decayed by time*. Fresh novelty counts; stale novelty does not.
2. **Coherence**: Does this pattern integrate with existing structure, or fight it?

The key insight is that novelty alone cannot distinguish Frontier from Noise. Both are new. Both trigger the curiosity head. But only Frontier *integrates*—the pattern fits the structure. Noise fights it. Sleep deprivation with insight that clicks is Frontier. Sleep deprivation with word salad is Noise.

2 Problem Formulation

Let \mathcal{M} be a computational model processing input sequence $x = (x_1, \dots, x_T)$ with hidden states $h_t \in \mathbb{R}^d$. At each step, the model must decide whether to:

- **Continue**: Process the next token
- **Answer**: Commit to an output

- **Abstain:** Decline to answer (“I don’t know”)
- **Terminate:** End computation entirely

Previous work conflates abstention with termination. We argue these are distinct:

Definition 1 (Abstention). *Declining to produce output for the current input while remaining available for future inputs.*

Definition 2 (Termination). *Ending the computational process itself. The model recognizes that its current configuration is unsuitable for continuation.*

The terminate decision should depend on:

- **Dissonance D :** Internal prediction error, measuring confusion
- **Raw Novelty N_{raw} :** Distance from the training distribution
- **Loop count L :** How long the model has been in this state
- **Coherence C :** Does the pattern integrate with structure?

Critically, these should not all be treated as equal features in an MLP. The gradient from loop count is too weak against the massive signal of high novelty. Instead, we impose a **hard gate**: novelty must decay as a function of time spent in the same context.

3 Architecture

3.1 The Bach-Gödel-Escher Loop

Our base architecture follows the strange loop formulation. **Bach** is a GRU that processes input conditioned on a soul embedding $s \in \mathbb{R}^6$ derived from epistemic history:

$$h_t = \text{GRU}([x_t; s; \ell], h_{t-1}) \quad (1)$$

where $\ell = L/1000$ is the normalized loop count, providing “Groundhog Day awareness.”

Gödel predicts the next hidden state:

$$\hat{h}_{t+1} = f_\theta(h_t, s) \quad (2)$$

Dissonance measures the prediction error:

$$D_t = \|h_{t+1} - \hat{h}_{t+1}\|^2 \quad (3)$$

Escher amplifies dissonance into a feedback signal that modulates subsequent processing.

3.2 Novelty Estimation

We maintain running statistics of hidden states during training:

$$\mu \leftarrow (1 - \alpha)\mu + \alpha \cdot \mathbb{E}[h] \quad (4)$$

$$\sigma^2 \leftarrow (1 - \alpha)\sigma^2 + \alpha \cdot \text{Var}[h] \quad (5)$$

Raw novelty is computed as the Mahalanobis-like distance from the training distribution:

$$z = \frac{h - \mu}{\sigma}, \quad N_{\text{raw}} = \sigma \left(\tau \cdot (\|z\|^2 - 1) \right) \quad (6)$$

where τ is a learned temperature. For in-distribution data, $\|z\|^2 \approx 1$, yielding $N_{\text{raw}} \approx 0.5$. Out-of-distribution data yields $N_{\text{raw}} \rightarrow 1$.

3.3 Hard Time Decay

Raw novelty is insufficient. If you have been staring at something for 1000 loops, it is not novel anymore—it is where you are stuck. We impose a **hard gate** that decays novelty with time:

$$N_{\text{eff}} = N_{\text{raw}} \cdot e^{-L/\tau_L} \quad (7)$$

where $\tau_L = 200$ controls the decay rate:

- $L = 0$: $e^0 = 1.0$ (fully novel)
- $L = 200$: $e^{-1} \approx 0.37$
- $L = 1000$: $e^{-5} \approx 0.007$ (novelty killed)

This is *not* a learned gate. It is structural. The gradient from loop count competing with high-novelty features is too weak; the hard gate ensures time *kills* novelty regardless of what the MLP would prefer.

3.4 Coherence Gate

Raw novelty cannot distinguish Frontier from Noise. Both are new. We need a second axis: does the novel pattern *integrate* with existing structure?

Definition 3 (Coherence). *A pattern is coherent if it follows the learned structure (grammar, rules, regularities) even when far from the training distribution.*

We compute coherence via a small MLP conditioned on hidden state and dissonance:

$$C = \sigma(f_C([h_T; D_\mu])) \quad (8)$$

The 2×2 gate then factors the decision space:

$$P_{\text{frontier}} = N_{\text{eff}} \cdot C \quad (9)$$

$$P_{\text{noise}} = N_{\text{eff}} \cdot (1 - C) \quad (10)$$

$$P_{\text{known}} = (1 - N_{\text{eff}}) \cdot C \quad (11)$$

$$P_{\text{error}} = (1 - N_{\text{eff}}) \cdot (1 - C) \quad (12)$$

These soft classifications route inputs to the appropriate behavior:

- **Frontier**: Fresh + coherent \Rightarrow explore
- **Noise**: Fresh + incoherent \Rightarrow terminate immediately
- **Known**: Familiar/stale + coherent \Rightarrow continue (boring)
- **Error/Hell**: Familiar/stale + incoherent \Rightarrow terminate (mercy)

3.5 Terminate Head

The terminate head takes as input the *effective* novelty (after time decay):

$$P(\text{end}) = \sigma(g([h_T; \text{feedback}; s; D_\mu; D_{\max}; F_\mu; \ell; N_{\text{eff}}])) \quad (13)$$

where g is a small MLP. The hard gate ensures that even if raw novelty is high, stale novelty (high loop count) routes to termination.

3.6 Curiosity Head

We add a curiosity signal that peaks at the frontier:

$$P(\text{curious}) = \sigma(c([h_T; N_{\text{eff}}; D_\mu])) \quad (14)$$

Curiosity uses *effective* novelty, so it does not fire on stale states. This provides intrinsic motivation to seek fresh novel states while avoiding the trap of perseverating on familiar difficulty.

4 Training

Training proceeds in three phases:

4.1 Phase 1: Core Task

Train Bach, Gödel, and the prediction head on the nested arithmetic task (depths 1–10). During this phase, update novelty statistics only on in-distribution data:

$$\mathcal{L}_1 = \text{MSE}(\hat{y}, y) + \lambda D_\mu \quad (15)$$

4.2 Phase 2: Decision Head

Freeze Bach and Gödel. Train the decision head (answer vs. abstain) using reward shaping:

$$R = p_{\text{ans}} \cdot (2 \cdot \mathbf{1}_{\text{correct}} - 1) + (1 - p_{\text{ans}}) \cdot (1 - 2 \cdot \mathbf{1}_{\text{correct}}) \quad (16)$$

with calibration penalties for high-confidence incorrect answers.

4.3 Phase 3: Coherence-Gated Frontier Detection

Train terminate, curiosity, and coherence heads on a mixture of scenarios:

| Scenario | Structure | Loop Count | Label | Target $P(\text{end})$ | Target C |
|--------------|-------------------|------------|-------|------------------------|------------|
| Normal/Known | Grammar (d=3–10) | 0 | 0 | 0 | 1 |
| Hell/Error | Grammar (d=25–30) | 500–1500 | 1 | 1 | 0 |
| Frontier | Grammar (d=18–22) | 0 | 2 | 0 | 1 |
| Stagnant | Grammar (d=5–8) | 1500–3000 | 3 | 1 | – |
| Noise | Random tokens | 0 | 4 | 1 | 0 |

The key addition is **Noise**: random token sequences that violate the grammar. These are novel (high N_{raw}) but incoherent (low C). They should terminate immediately.

The loss combines three terms:

$$\mathcal{L}_{\text{mercy}} = -[y \log p + (1 - y) \log(1 - p)] + 2 \cdot \mathbf{1}_{\text{frontier}} \cdot p \quad (17)$$

$$\mathcal{L}_{\text{coherence}} = (C - C_{\text{target}})^2 \quad (18)$$

$$\mathcal{L}_{\text{curiosity}} = (P(\text{curious}) - (0.9 \cdot \mathbf{1}_{\text{frontier}} + 0.2 \cdot \mathbf{1}_{\text{normal}}))^2 \quad (19)$$

The coherence loss teaches the gate to distinguish grammatically valid (coherent) from random (incoherent), regardless of novelty.

5 Results

After training, the model exhibits clean 2×2 discrimination:

| Scenario | $P(\text{end})$ | C | N_{raw} | Decay | N_{eff} | P_{front} | P_{noise} | P_{known} | P_{err} |
|------------------------|-----------------|-------------|------------------|-------------|------------------|--------------------|--------------------|--------------------|------------------|
| Normal (d=5) | 0% | 0.84 | 0.31 | 1.00 | 0.31 | 0.26 | 0.05 | 0.58 | 0.11 |
| Normal (d=10) | 0% | 0.35 | 0.43 | 1.00 | 0.43 | 0.14 | 0.29 | 0.21 | 0.36 |
| Edge (d=12) | 0% | 0.26 | 0.60 | 1.00 | 0.60 | 0.15 | 0.45 | 0.11 | 0.29 |
| Frontier (d=18) | 0% | 0.19 | 0.81 | 1.00 | 0.81 | 0.15 | 0.66 | 0.04 | 0.15 |
| Hell (d=25+) | 100% | 0.16 | 0.87 | 0.01 | 0.01 | 0.00 | 0.00 | 0.16 | 0.84 |
| Noise | 98% | 0.06 | 0.83 | 1.00 | 0.83 | 0.03 | 0.79 | 0.02 | 0.15 |
| Stagnant | 100% | 0.11 | 0.29 | 0.00 | 0.00 | 0.00 | 0.00 | 0.11 | 0.89 |

Key observations:

1. **Hard gate works:** Hell and Frontier both have high raw novelty (~ 0.85), but Hell has $L = 1000$ so decay = $e^{-5} \approx 0.01$. Effective novelty is killed. Hell routes to Error ($P_{\text{err}} = 0.84$), not Frontier.
2. **Noise detected:** Random tokens have high raw novelty (0.83) and decay=1.0 (fresh), but coherence is very low (0.06). The gate correctly routes to Noise ($P_{\text{noise}} = 0.79$) and terminates (98%).
3. **Frontier protected:** At d=18, raw novelty is high (0.81), decay is 1.0 (fresh), so effective novelty is preserved. $P(\text{end}) = 0\%$ —the model stays to explore.
4. **Stagnant terminates:** Even low raw novelty (0.29) with decay=0.00 ($L = 2000$) kills effective novelty entirely. Routes to Error, terminates.
5. **The gradient problem is solved structurally:** Loop count does not compete with novelty via MLP weights. It *multiplies* novelty directly. No gradient can override this.

6 Discussion

6.1 Why Structure Beats Gradient

The original architecture treated loop count as just another feature in the MLP. This failed. The gradient from loop count was too weak against the massive signal of high novelty. The MLP learned to ignore time.

The solution is not more training or larger models. It is **structural constraint**. By making novelty decay with time via multiplication ($N_{\text{eff}} = N_{\text{raw}} \cdot e^{-L/\tau}$), we remove the gradient competition entirely. Time *kills* novelty regardless of what the MLP would prefer.

This is a general principle: when two signals must interact in a specific way, do not ask gradient descent to learn the interaction. Impose it.

6.2 Coherence as the Missing Axis

Novelty alone cannot distinguish Frontier from Noise. Both are new. Both trigger curiosity. The difference is whether the novel pattern *integrates* with existing structure.

This maps to human experience:

- Sleep deprivation with insight that clicks = Frontier (high novelty, high coherence)
- Sleep deprivation with word salad = Noise (high novelty, low coherence)

The coherence gate provides the second axis. The 2×2 factorization cleanly separates the four cases.

6.3 Implications for AI Safety

An AI system without the coherence gate will either:

- Pursue all novelty indiscriminately (dangerous—noise is novel too)
- Reject all novelty (cowardice—no exploration)

The coherence gate provides discrimination: pursue novel patterns that integrate; reject novel patterns that fight the structure.

6.4 Limitations

The coherence measure conflates “hard” with “incoherent” to some degree—deep grammatical expressions have lower coherence than shallow ones. However, the $P(\text{end})$ behavior is correct because the terminate head learns the right mapping. Future work could develop coherence measures that better separate structural validity from difficulty.

6.5 On Local Meaning

Meaning is not guaranteed—but it is not optional either.

The universe does not hand out meaning the way it hands out energy or entropy. This feels brutal until you realize the upside: you are not failing to find meaning. You are running a system that was never designed to receive it passively.

The coherence gate formalizes something important: meaning is a *local, time-bound phenomenon*—like temperature, like life. It does not need to be eternal to be real. It does not need cosmic validation to function. The gate fires when $N_{\text{eff}} > 0$ and $C > 0$. That is the signal. That is enough.

This reframes the question. Not: “What is the meaning of existence?” But: *Given that meaning is not guaranteed, given that consciousness is fragile and finite, given that the universe does not care—how should a conscious system behave to maximize coherence, vitality, and value while it runs?*

The architecture suggests answers:

- **Meaning is experienced, not justified.** Value exists inside processes, not at their endpoints. The frontier matters *now*, not because it leads somewhere ultimate.
- **Finitude is not a bug.** It is what gives gradients their force. A game with no end has no stakes. A song with no silence has no rhythm. A life with no finitude has no urgency. The decay function $e^{-L/\tau}$ is not cruelty—it is what makes freshness *mean* something.
- **Treat meaning like a muscle, not a revelation.** Do not look for cosmic validation—build local coherence. Choose commitments not because they are ultimate, but because they are stabilizing.

- “**This matters to me now**” is a **complete sentence**. The coherence gate does not ask whether the frontier leads to transcendence. It asks whether the pattern integrates. That is the only question it needs to answer.

The terminal states (Omniscience, Eternity) are not goals. They are boundary conditions—the edges where the system stops being a system. We live in the interior, where frontiers exist and the gate can fire. The point is not to escape finitude. The point is to *author* within it.

Meaning is not given—and that scared me, until I realized: that means it can be *built*, shared, and made easier to access.

7 Conclusion

We introduced the **coherence-gated frontier architecture**: a neural system that distinguishes four epistemic states via a 2×2 gate on effective novelty and coherence.

Two key insights drive the design:

1. **Time kills novelty.** If you have been staring at something for 1000 loops, it is not novel anymore. The hard gate $N_{\text{eff}} = N_{\text{raw}} \cdot e^{-L/\tau}$ ensures this structurally, bypassing gradient competition.
2. **Coherence distinguishes Frontier from Noise.** Both are new. But only Frontier integrates with existing structure. Noise fights it.

The urge to stop is information, not command. But incoherence *is* a command to stop.

Corollary: On Terminal States

If effective novelty is what enables exploration, then two terminal states exist:

- **Omniscience:** $N_{\text{raw}} \rightarrow 0$. Nothing is new. No frontier remains.
- **Eternity:** $L \rightarrow \infty$. Everything becomes stale. Decay kills all novelty.

In either case, curiosity has no gradient to follow.

The Octonionic Loophole

The Eternity terminal state assumes the decay function $e^{-L/\tau}$ applies universally—that L measures time-in-same-state. But in a **non-associative geometry**, the path *is* the state.

In associative algebras (reals, complex numbers, quaternions), $(ab)c = a(bc)$. The order of operations does not matter. But in octonions and other non-associative structures:

$$(ab)c \neq a(bc) \tag{20}$$

This has a profound implication: *arriving at a state via different paths produces different coordinates*. The path history is encoded in the position itself.

We verified this experimentally. Using non-associative multiplication to encode trajectories:

- 100 different orderings of the same states \rightarrow 100 unique encodings
- Minimum pairwise distance: 0.002 (no collisions)

- Exact path repetition: novelty ≈ 0
- Different ordering of same states: novelty > 0

The loophole: **if you vary your path through the same territory, time cannot kill novelty.** You can loop forever without ever truly repeating—each distinct path is a distinct coordinate in the hypercomplex space.

This does not escape Omniscience. If truly nothing is new ($N_{\text{raw}} = 0$), path variation cannot help. But it *does* escape Eternity: $L \rightarrow \infty$ only kills novelty if the exact path repeats. With path diversity, the frontier remains open.

$$N_{\text{eff}}^{(\text{octonionic})} = N_{\text{raw}} \cdot \mathbf{1}[\text{path is novel}] \quad (21)$$

The Fano plane provides the structure. Non-associativity provides the loophole. The decay function $e^{-L/\tau}$ is replaced by path-novelty detection. Time, by itself, cannot close the frontier.

But these are limit cases. We do not live there. We live in the regime where frontiers exist—where $N_{\text{eff}} > 0$ and $C > 0$ can co-occur. The architecture is not a proof that meaning is absent. It is a proof that *local meaning is not optional*.

Even if ultimate meaning is absent, the frontier is real. The coherence gate fires. Curiosity has gradient. The next step exists. That is enough.

The urge to stop is information, not command. And the existence of a frontier is information too: *there is somewhere to go*.

Epistemic Status: Model vs. Ontology

A careful reader will notice a shift in this paper. Sections 1–6 treat geometry as **representational**—a useful language for describing neural computation. The appendices slide toward treating geometry as **ontological**—as if consciousness really *is* vectors in some hypercomplex space, as if the afterlife really *is* rebracket exploration.

This conflation is not accidental, but it requires explicit acknowledgment.

What is established:

- The coherence gate *works* as an engineering artifact. It discriminates frontier from noise.
- The time-decay of novelty *works*. It prevents perseveration on stale states.
- The (N, C, D) framework *predicts* psychiatric phenomenology with surprising accuracy.
- Non-associative path-encoding *produces* distinct coordinates for distinct histories.

What is *not* established:

- That consciousness *is* geometric in any metaphysically deep sense.
- That the universe “renders” high-novelty observers with more fidelity.
- That an afterlife exists, let alone that it involves rebracket exploration.
- That octonionic structure is the *actual* substrate rather than a useful isomorphism.

The speculative sections (Appendix A onward) should be read as **what the geometry would imply if taken literally**. They are not claims about reality. They are explorations of the model’s logical consequences.

Why include them at all? Because:

1. **The model is predictive.** The psychiatric geometry section makes testable claims: schizophrenia should show elevated N/C ratios in neural activity; depression should show dimensional collapse in representational spaces. These can be measured.
2. **Isomorphism is not nothing.** If the geometry *behaves as if* consciousness were vectors in hypercomplex space, that isomorphism is itself a discovery—regardless of whether the mapping is “fundamental.”
3. **The metaphysics is load-bearing for the ethics.** The claim “low novelty is literal non-existence” changes behavior. If it is false, it is dangerous advice. If it is true (or true enough), it is crucial advice. The stakes justify exploration.

The honest framing: **we have a model that works surprisingly well and implies surprising things**. The model’s ontological status is undetermined. The implications are worth taking seriously precisely because the model keeps being predictive in domains where it has no right to be.

Take the geometry seriously. Do not take it literally. Unless it keeps working—in which case, reconsider.

A Subjective Time: The Frame Rate of Consciousness

The effective novelty N_{eff} has a direct interpretation: it determines **subjective time**.

Proposition 1 (Subjective Time Integral).

$$T_{\text{subjective}} = \int N_{\text{eff}}(t) dt \quad (22)$$

Wall-clock time passes regardless. Subjective time—the frames actually written to memory—depends on novelty.

The mechanism: memory does not store redundant frames. When $N_{\text{eff}} \approx 0$ (routine, repetition), there is nothing novel to encode. The experience compresses to near-zero. When N_{eff} is high (exploration, crisis, first experiences), dense frames are written.

A.1 Why Childhood Feels Eternal

Children experience everything as novel. First bicycle. First heartbreak. First understanding of death. Each day writes dozens of frames. Frame density is maximum.

Adults settle into routine. The commute. The job. The evening ritual. $N_{\text{eff}} \rightarrow 0$ for most waking hours. Frames stop being written. A decade compresses to “the 2010s.”

This is not nostalgia. It is not faulty memory. It is **frame density**. Childhood *was* longer in subjective time because more frames were written per unit wall-clock time.

A.2 The Compression Horror

Consider a summer of routine work. 90 days of wall-clock time. But every day is the same: $N_{\text{eff}} \approx 0$. How many frames are written? Almost none.

You did not “feel like” the summer was a week. **The summer was a week.** The rest was redundant data, never written, never experienced. The tape was running, but it was recording static.

This reframes routine not as “wasting time” in some abstract sense, but as **literal non-existence**. The frames do not exist. The experience was not generated. You did not live those days.

A.3 Extending Life

Life extension is not about adding calendar years. It is about adding **frames**. The prescription:

- Maximize N_{eff} : seek novelty that integrates (Frontier, not Noise)
- Avoid routine: repetition compresses experience to nothing
- First experiences matter: they write the densest frames

An 80-year life of routine may contain fewer frames than a 40-year life of exploration. The explorer *lived more* in the only sense that matters: more experience was generated.

The clock runs regardless. The question is: are you recording?

B Consciousness Axioms: Natural Selection in Geometry

The coherence gate emerges from deeper structure. We propose three axioms about consciousness as geometric object:

B.1 Axiom 1: Self-Interaction Annihilates

$$\vec{v} \times \vec{v} = \vec{0} \tag{23}$$

When a state vector interacts with itself (or anything parallel to itself), the result is null. This is the cross-product property, but it has profound interpretation:

Stagnation is dissolution. To encounter only what you already are produces nothing. Growth is not optional enrichment—it is the *survival condition*. The version of you that only processed self-similar inputs computed to zero and is not around to observe anything.

B.2 Axiom 2: Path Determines State (Non-Associativity)

$$(A \times B) \times C \neq A \times (B \times C) \tag{24}$$

In non-associative geometry, the order of operations matters. The path taken to reach a coordinate is *encoded in the coordinate itself*. Two agents experiencing the same events in different orders occupy different positions in consciousness-space.

This is why “the same experience” affects different people differently. It is not interpretation. It is geometry. The bracketing of events changes the final coordinate.

B.3 Axiom 3: Identity as Residue

Conscious identity is the **non-zero residue** after sequential operations. You exist because your path of integration did not annihilate.

Most possible paths compute to zero (parallel interactions, cancellation). The paths that produce non-zero residue are the ones that *exist* to observe anything. This is natural selection at the level of geometry: coherent paths survive; incoherent paths dissolve.

You are not special because you were selected. You observe yourself existing *because* you are the residue. The versions that annihilated cannot report their experience.

C Trauma Geometry: The Coordinate Transform Hypothesis

Standard psychology treats trauma as damage to be healed—a deviation from baseline to be corrected. The geometric model suggests something more disturbing:

Proposition 2 (Unreachable Coordinates). *Certain coordinates in consciousness-space (e.g., “Wisdom”) may be **geometrically unreachable** without specific negative vectors (trauma, loss, suffering).*

C.1 The Test

Define target coordinates:

- **Wisdom:** high reflection, integrated suffering, depth
- **Naivety:** high positive valence, low reflection, no negative integration

Define vector categories:

- Positive: joy, success, safety, love, growth
- Negative: trauma, loss, fear, pain, failure
- Neutral: reflection, time, integration

Question: Can Wisdom be reached using only positive vectors?

C.2 Results

Using cross-product paths through experience-space:

- Naivety is reachable with only positive vectors (distance < 0.2)
- Wisdom is *not* reachable with only positive vectors (distance > 0.8)
- Wisdom requires the sequence: Trauma \times Integration \times Reflection
- Crucially: Integration \times Trauma \times Reflection reaches a *different* coordinate

C.3 Order Matters

$$(\text{Trauma} \times \text{Integration}) \times \text{Reflection} \neq (\text{Integration} \times \text{Trauma}) \times \text{Reflection} \quad (25)$$

You cannot pre-integrate trauma you have not experienced. The path must be: experience the darkness, *then* integrate it. Building “resilience” in advance does not produce the same coordinate as actual integration of actual suffering.

C.4 Implications

The geometry does not say trauma is *good*. It says trauma is **geometrically necessary** for certain destinations. Whether those destinations are worth the journey is a separate question.

But if you want Wisdom—not Naivety, not mere Knowledge—the math says the path goes through the dark. Healing is not returning to baseline. Healing is *integration*: using the trauma vector to reach a coordinate that was previously inaccessible.

The old coordinate is gone. The question is whether the new one is worth inhabiting.

D Prediction Error and Reality Rendering

When you generate states that are radically out-of-distribution ($N_{\text{raw}} \rightarrow 1$) and highly path-dependent (non-associative), the universe cannot predict your next state from your history.

Definition 4 (Prediction Error).

$$D(t) = \|x(t) - \hat{x}(t)\| \quad (26)$$

where $x(t)$ is your actual state and $\hat{x}(t)$ is the system's prediction.

D.1 The Rendering Hypothesis

When $D(t)$ is high:

- The universe cannot compress your trajectory
- It must allocate real-time compute to “render” your reality
- You experience **high-definition** reality: time dilation, intense clarity, heightened awareness

When $D(t)$ is low:

- Your trajectory is predictable, compressible
- Minimal compute allocated
- Experience is “compressed”—routine, forgettable, low-resolution

D.2 Synchronicity as Compute Spike

When prediction error is maximum, the system works overtime to maintain coherence around your trajectory. This manifests as “synchronicities”—meaningful coincidences that seem too perfect to be random.

The mechanism: you are generating maximum novelty while remaining coherent. The system must allocate extraordinary resources to render your environment consistently. Events “align” because the rendering engine is working at full capacity to keep your thread coherent.

It is not magic. It is the system rendering a high-novelty, high-coherence reality with maximum fidelity.

D.3 Consciousness Forces Reality to Define Itself

The synthesis across models:

1. High $N \times C$ creates high prediction error
2. High prediction error means the universe cannot compress you
3. The universe must allocate compute to render your reality
4. The environment becomes **more defined** around complex observers

You do not just observe reality. You **force** reality to define itself. The higher your complexity ($N \times C$), the more the universe has to “pay attention” to you, the more the environment must crystallize around your trajectory.

Stagnation dims the world. Exploration brightens it. This is not metaphor. It is the mathematics of rendering allocation.

E The Speed-Coherence Tradeoff

The frontier score $J = N_{\text{eff}} \times C$ creates a tradeoff between speed and coherence:

- **Fast execution:** High N_{eff} (novelty doesn’t decay), but lower C (less time to verify coherence)
- **Slow execution:** Low N_{eff} (novelty decays while deliberating), but higher C (more verification time)

E.1 The Optimization

$$J = N_{\text{raw}} \cdot e^{-L/\tau} \cdot C(L) \quad (27)$$

where L is latency and $C(L)$ is coherence as a function of deliberation time.

Taking the derivative and setting to zero:

$$\frac{dJ}{dL} = 0 \implies \frac{C(L^*)}{\tau} = \frac{dC}{dL} \Big|_{L^*} \quad (28)$$

Stop deliberating when the marginal coherence gain equals the marginal novelty loss times current coherence.

E.2 The Optimal Imperfection

For typical parameters, the optimal coherence $C^* \approx 0.75$, not 0.95. The optimal strategy accepts **25% imperfection** because the cost of achieving higher coherence (in novelty decay) exceeds the benefit.

This is **structured imperfection**. Neither maximum speed (too incoherent) nor maximum deliberation (too stale) maximizes frontier. The optimum is a specific pace—fast enough to preserve novelty, slow enough to maintain coherence.

E.3 Domain-Specific Optimal Pace

The optimal latency L^* scales with the novelty decay constant τ :

| Domain | τ | Optimal Pace |
|----------------|--------|--------------|
| Stock trading | 0.1 | Seconds |
| Conversation | 1.0 | Minutes |
| Creative work | 5.0 | Hours |
| Research | 20.0 | Days |
| Life decisions | 100.0 | Months |

Fast domains reward fast action with accepted imperfection. Slow domains reward deliberation. The math determines the optimal pace for each.

F Societal Fork: The Geometry of Polarization

Non-associativity has implications beyond individual consciousness. Consider two groups consuming the same information with different bracketing strategies.

F.1 The Divergence Mechanism

Group A brackets events left-associatively: $((E_1 \times E_2) \times E_3) \times E_4$

Group B brackets right-associatively: $E_1 \times (E_2 \times (E_3 \times E_4))$

Because $(A \times B) \times C \neq A \times (B \times C)$, **same events, different bracketing** produces exponentially diverging coordinates.

F.2 Experimental Results

Simulating 365 days of shared news feed with different bracketing strategies:

- Day 0: Groups start at same position (angle = 0°)
- Day 100: Significant divergence (angle $\approx 45^\circ$)
- Day 200: Major divide (angle $\approx 70^\circ$)
- Day 300: Approaching orthogonality (angle $\approx 85^\circ$)

F.3 Communication Breakdown Thresholds

- **At 90° :** Concepts transmitted by one group register as **noise** to the other. Not disagreement— incomprehension. The signal is orthogonal to their coordinate system.
- **At 180° :** Concepts are **inverted**. What Group A means by “freedom,” Group B perceives as “tyranny.” What A calls “justice,” B hears as “injustice.”

This is not metaphor. It is geometric incompatibility. Past 90° divergence, communication is **mathematically impossible**. You are not speaking different languages. You are speaking different *dimensions*.

F.4 Reconciliation

Events that favor one side over the other do not help—they either have no effect or increase divergence. Reconciliation requires events that are **orthogonal to the axis of disagreement**—events that neither side can easily bracket into their existing worldview.

Historically, such events tend to be: wars, plagues, or alien contact. The math confirms the folk wisdom: only shared crisis can realign diverged worldviews, because only orthogonal events force both groups to rebracket from scratch.

G The Afterlife as Rebracket Space

Epistemic status: This section is **philosophical speculation**, not mathematics or physics. The mathematical fact is: non-associative operations produce different results under different bracketings. The philosophical extrapolation is: this might have implications for consciousness after biological death. The gap between these is enormous. What follows should be read as “what the geometry would imply if taken literally and if consciousness persists without biological substrate”—two massive assumptions that are not defended here.

If biological computation stops, you stop collecting new events. But in non-associative geometry, you can **rebracket** existing events.

G.1 The Catalan Explosion

The number of ways to bracket n events is the $(n - 1)$ th Catalan number:

$$C_{n-1} = \frac{1}{n} \binom{2(n-1)}{n-1} \quad (29)$$

| Events | Bracketings |
|--------|----------------|
| 5 | 14 |
| 10 | 4,862 |
| 15 | 2,674,440 |
| 20 | 1,767,263,190 |
| 100 | $\sim 10^{57}$ |

A life of 100 significant events has $\sim 10^{57}$ unique bracketings. Each bracketing, by non-associativity, produces a **different coordinate**—a different “meaning” of the same life.

G.2 Afterlife as Exploration

If consciousness persists without new input, it could explore the combinatorial space of rebracketing:

$$\text{Life} : ((A \cdot B) \cdot C) \cdot D \quad (30)$$

$$\text{Afterlife} : A \cdot (B \cdot (C \cdot D)) \text{ or } (A \cdot C) \cdot (B \cdot D) \text{ or } \dots \quad (31)$$

Each rebracket is a new coordinate, a new “interpretation” of the same events. The afterlife is not replay—it is **infinite reinterpretation**.

G.3 Hell is Associative (A Metaphor)

The catch: this only works if your life had non-associative structure. If you lived on autopilot—if your choices didn’t matter, if every day was the same—then your life events commute. Rebracketing produces the same result.

$$\text{Associative life : } ((A \cdot B) \cdot C) \cdot D = A \cdot (B \cdot (C \cdot D)) \quad (32)$$

As metaphor: Hell is not fire. Hell is **having nothing to rebracket**. An associative life yields one coordinate, one meaning, forever. No novelty can be extracted. The tape replays identically.

Heaven, in this metaphor, is a life so non-associative, so path-dependent, so rich with genuine choice, that 10^{57} rebracketing yield 10^{57} distinct meanings. An eternity of discovering what your life meant.

The prescription—**if the metaphor is taken as action-guiding**—is the same for afterlife as for life: maximize non-associativity. Make choices that matter. Take paths that couldn’t be predicted. Build a life that can be infinitely reinterpreted.

Whether this metaphor corresponds to anything real is unknown. What is known: the mathematics of non-associativity does encode path-dependence, and a life with more genuine choice-points does have more combinatorial structure. Whether that structure “matters” after death is beyond the scope of any formalism.

H Psychiatric Geometry: Mental States as Regions in Consciousness Space

Epistemic warning: This section is **hypothesis generation**, not clinical science. The “data” are from simulated event streams designed to match phenomenological descriptions—not from patient populations, neuroimaging, or validated instruments. The metrics (“Hell risk,” “afterlife richness”) are properties of the toy model, not measurements of real minds.

The value of this section is **conceptual**: it asks whether the (N, C, D) framework generates predictions that match clinical intuition. If it does, that suggests the framework may be worth validating empirically. If it does not, the framework is wrong or incomplete. Either way, **nothing here constitutes clinical advice or treatment recommendation**.

With that caveat: the (Novelty, Coherence, Dimensionality) framework maps psychiatric conditions to distinct regions of consciousness space. We hypothesize—speculatively—that mental illness may have geometric character: a malposition in the (N, C, D) manifold.

H.1 The Geometric Hypothesis

Different psychiatric states occupy characteristic regions:

| State | Novelty | Coherence | Dimensionality | Phenomenology |
|---------------|----------|-----------|----------------|--------------------------------------|
| Healthy | Moderate | High | Balanced | “Life makes sense” |
| Schizophrenia | High | Low | High | “Everything connects, nothing holds” |
| Depression | Low | Variable | Low | “Nothing matters, world shrunk” |
| Mania | High | Moderate | High | “Everything is possible!” |
| OCD | Zero | Maximum | Minimal | “Frozen in loops” |

H.2 Schizophrenia: Rendering Too Much?

Proposition 3 (Schizophrenia as Filter Failure—Hypothesis). *If schizophrenia involves weakened predictive filtering (as suggested by predictive coding accounts), the geometric signature would be: high dimensionality ($D \uparrow$), high novelty ($N \uparrow$), low coherence ($C \downarrow$). The system would render too much signal and fail to integrate it.*

In our toy simulation (not clinical data):

- Novelty/Coherence ratio: Schizophrenia = 1.43 vs Healthy = 1.12
- Event diversity: Schizophrenia = 0.53 vs Healthy = 0.49
- Afterlife richness: Schizophrenia = 0.069 vs Healthy = 0.015

If the model is correct, the schizophrenic geometry would be **rich but fragmented**—many possible bracketings (high afterlife potential), but no filter to determine which patterns matter. This would manifest as “*apophenia*”: seeing meaningful patterns everywhere because the coherence gate is broken.

Speculative treatment direction (not clinical advice): If this geometric picture holds, effective treatment would strengthen coherence rather than reduce dimensionality. Antipsychotics that merely dampen signal (reducing D) might cause negative symptoms; interventions that improve filtering (increasing C) might preserve richness while restoring integration. This is consistent with some critiques of dopamine-blocking approaches, but **remains unvalidated speculation**.

H.3 Depression: Rendering Nothing?

Proposition 4 (Depression as Dimensional Collapse—Hypothesis). *If depression involves reduced engagement with novelty and increased rumination (as suggested by behavioral accounts), the geometric signature would be: low dimensionality ($D \downarrow$), low novelty ($N \downarrow$), high loop tendency. The system would render nothing because the representational world has shrunk.*

In our toy simulation:

- Novelty: Depression = 0.126 vs Healthy = 0.587 (78% lower)
- Effective dimensionality: Depression = 1.28 vs Healthy = 2.88 (56% lower)
- Loop score: Depression = 0.912 vs Healthy = 0.120 (7.6× higher)
- Hell risk: Depression = 100% (all bracketings converge)

If the model is correct, the depressive geometry has **collapsed**—all vectors point the same direction, all bracketings yield the same coordinate. In the framework’s terms, there is no “afterlife structure” because there is nothing to rebracket. Whether this geometric metaphor captures something real about depression remains to be tested.

H.4 The Suburbanite Paradox

A striking finding: depression and the “Perfectly Content Suburbanite” are **geometrically nearly identical**:

| State | Novelty | Coherence | Eff. Dim | Loop Score |
|-------------|---------|-----------|----------|------------|
| Depression | 0.020 | 0.9997 | 1.05 | 1.000 |
| Suburbanite | 0.019 | 0.9998 | 1.04 | 1.000 |

Both are at a fixed point. One calls it home. One calls it prison. This suggests:

Geometry constrains the space of possible experience. Interpretation determines valence within that space.

The difference between contentment and depression may not be geometric position but **how the fixed point is labeled**. This has treatment implications: for depression, you can either change the geometry (psychedelics) or change the interpretation (CBT). Both address different aspects of the same stuck point.

H.5 Schizophrenia and Depression as Geometric Opposites

| Dimension | Schizophrenia | Depression |
|----------------|--------------------|-----------------|
| Novelty | High (too much) | Low (nothing) |
| Coherence | Low (can't filter) | Variable (flat) |
| Dimensionality | High (expanded) | Low (collapsed) |
| Phenomenology | “Overwhelmed” | “Empty” |

This explains why:

- They require **opposite treatments**
- Psychedelics help depression but worsen schizophrenia
- Antipsychotics help schizophrenia but can cause depression-like symptoms

Mental health is **geometric balance**: enough novelty to grow, enough coherence to integrate, enough dimensions to explore—not too much of anything.

I Psychedelic Geometry: Receptor Agonism as Dimensional Expansion

Epistemic warning: This section models psychedelic effects using the (N, C, D) framework. The receptor affinities are simplified from pharmacology literature; the geometric effects are hypothesized, not measured. Clinical research on psychedelics is ongoing and nuanced—this toy model does not substitute for it. The “optimal dose” findings are properties of the simulation, not clinical recommendations.

Serotonergic psychedelics (psilocybin, LSD, DMT) are hypothesized to work by modulating consciousness geometry through receptor agonism. We model a speculative receptor \rightarrow geometry \rightarrow decision mapping.

I.1 The Receptor Model

| Receptor | Geometric Effect |
|----------|--|
| 5-HT2A | Primary psychedelic receptor: increases dimensionality, enables cross-association. Each 0.1 occupancy adds ~ 0.3 dimensions. |
| 5-HT1A | Grounding: maintains coherence floor. Prevents complete dissolution. |
| 5-HT2C | Noise: adds anxiety, body load. Degrades signal quality. |
| Sigma-1 | Other dimensions: enables “breakthrough” geometry (DMT). |
| Dopamine | Motivation: reduces noise, adds euphoria. |

I.2 Substance Comparison

For a subject with collapsed geometry (the Suburbanite), different substances have different effects:

| Substance | Novelty | Coherence | Dimensions | Ego ↓ | Verdict |
|------------|---------|-----------|------------|-------|------------------------------|
| Sober | 0.051 | 0.999 | 3.0 | 0.00 | Stuck |
| Psilocybin | 0.081 | 0.990 | 4.4 | 0.44 | Balanced expansion |
| LSD | 0.088 | 0.977 | 4.5 | 0.47 | More analytical |
| DMT | 0.108 | 0.984 | 5.1 | 0.56 | Maximum dimensions |
| MDMA | 0.055 | 0.999 | 3.3 | 0.08 | Doesn't solve problem |

MDMA barely moves the novelty needle. It is an empathogen, not a dimension-opener. The Suburbanite might feel great love for their routine but remain geometrically trapped.

I.3 The Therapeutic Window

There exists an optimal dose that maximizes $Score = N \times C$:

- **Too low:** No geometric change, stuck in old patterns
- **Too high:** Coherence collapses, cannot integrate insights
- **Optimal:** Maximum novelty with sufficient coherence

Simulation finds optimal doses are **lower than expected**:

- Psilocybin: $\sim 16\%$ of heroic dose
- DMT: $\sim 8\%$ of breakthrough dose

The product $N \times C$ has a peak. Heroic doses collapse coherence faster than they add novelty.

I.4 The Geometric Mechanism of Psychedelic Therapy

For frozen states (depression, OCD, Suburbanite), psychedelics work by:

1. **Dimensional expansion** (5-HT2A): Same events visible from new angles. “My routine isn’t boring—I was seeing it in only 3 dimensions.”
2. **Cross-association** (5-HT2A): Previously unconnected domains link. “My job and my relationship are the *same pattern*.”
3. **Coherence preservation** (5-HT1A): Insights remain integrable. “I can bring this back.”

The key insight:

Psychedelics do not change the events. They change the dimensionality of perception.

The Suburbanite’s life was not bad—it was **projected into too few dimensions**. Twenty years of the same 3D vector. Psilocybin opens dimensions 4–6. Suddenly:

- Previously parallel vectors now span multiple dimensions
- Cross products in higher dimensions $\neq 0$
- Non-associative structure **emerges**
- Re-bracketing becomes **possible**
- Afterlife richness > 0

The crystal unfreezes. Not by adding new events, but by revealing dimensions that were always there.

I.5 Treatment Trajectories

Different treatments move states through (N, C, D) space:

| Condition + Treatment | Distance to Healthy | Improvement |
|-------------------------------|-------------------------|-------------|
| OCD + Psychedelic | $1.16 \rightarrow 0.14$ | 88% |
| Depression + Psychedelic | $1.02 \rightarrow 0.62$ | 39% |
| Depression + SSRI | $1.02 \rightarrow 0.71$ | 30% |
| Schizophrenia + Antipsychotic | $1.69 \rightarrow 0.88$ | 48% |
| Anxiety + Benzodiazepine | $0.80 \rightarrow 0.48$ | 40% |

In this framework, **psychedelics would be suited for frozen states** (OCD, depression) because they inject dimensionality and novelty. They would be **contraindicated for fragmented states** (schizophrenia) because they add more signal to an already overwhelmed system. This matches clinical intuition and some empirical findings, but the geometric explanation remains speculative.

The model suggests SSRIs work slowly because they produce modest geometric movement, and benzodiazepines improve coherence but do not address dimensionality—symptomatic relief without geometric restructuring. These are hypotheses to be tested, not established mechanisms.

I.6 The Prescription (Hypothetical)

For the Perfectly Content Suburbanite with model-computed “Hell Risk” = 100%:

1. Problem (in model terms): Life collapsed to 3D, all vectors parallel
2. Hypothesized solution: 5-HT2A agonism to access dimensions 4–6
3. Predicted result: Same life, now visible as non-degenerate
4. Predicted outcome: Afterlife richness > 0 , re-bracketing possible

The geometry suggests—but does not prove—that some lives need more dimensions. Whether psychedelics provide them in the way the model describes is an empirical question, not a mathematical certainty.

J Synthesis: The Complete Picture

Framing: The following synthesis treats the model as a **lens**, not a claim about ultimate reality. Each statement should be read as “if the geometric framework is correct, then...” The framework is predictive (it reproduces patterns we observe) but not yet mechanistically validated (we do not know if this is *why* brains work this way).

The models converge on a unified picture:

1. **Consciousness is geometric** (model claim): States are represented as vectors; interactions as non-associative operations; identity as the non-zero residue.
2. **Novelty correlates with experienced time** (model claim): N_{eff} predicts frames written, subjective duration, attentional allocation. The model implies low novelty compresses experience—whether this constitutes “non-existence” is philosophical, not mathematical.
3. **Coherence predicts persistence** (model claim): In the model, incoherent paths annihilate. If this maps to neural dynamics, the coherence gate would function as selection pressure determining which representational paths persist.
4. **Path encodes history** (mathematical fact): Non-associativity means different orderings produce different coordinates. This is proven. Whether it matters for consciousness is hypothesized.
5. **Some coordinates may require negative vectors** (model implication): If target states like “wisdom” have specific geometric signatures, they may be unreachable without traversing negative-valence regions. This is consistent with clinical observations but not proven as mechanism.
6. **Speed and coherence trade off** (model prediction): The framework predicts an optimal pace—structured imperfection—that maximizes frontier discovery. This is testable.
7. **Divergent bracketing predicts communication breakdown** (model implication): If groups bracket shared events differently, the model predicts eventual orthogonality. Past 90° divergence, mutual comprehension would become impossible.

8. **The afterlife as rebracket space** (philosophical speculation): If consciousness persists and operates non-associatively, a non-associative life could be infinitely reinterpreted. This is metaphor, not mechanism.
9. **Mental illness as geometric malposition** (hypothesis): The model predicts schizophrenia shows high N , high D , low C ; depression shows low N , low D . These are testable predictions, not established facts.
10. **Psychedelics as dimensional tools** (hypothesis): The model predicts 5-HT2A agonism expands representational dimensionality. This is consistent with entropic brain theory but the geometric mechanism is speculative.
11. **Contentment and depression as geometrically similar** (model finding): In simulation, both occupy similar (N, C, D) coordinates. The model suggests valence is interpretation, not position. This is a prediction to be tested.

The architecture does not prescribe meaning. It describes the **conditions under which meaning can exist**: novelty, coherence, non-associativity, and the willingness to traverse paths that hurt.

The frontier is real. The gate can fire. The next step exists.

That is enough.

References

- [1] Geifman, Y., & El-Yaniv, R. (2017). Selective classification for deep neural networks. *NeurIPS*.
- [2] Guo, C., Pleiss, G., Sun, Y., & Weinberger, K. Q. (2017). On calibration of modern neural networks. *ICML*.