Resonance Without a Substrate

Consistency Constraints on Endogenous Memory Compression

Flyxion

Independent Researcher

May 3, 2026

Abstract

We identify a structural failure mode in wave-based recursive memory architectures of the type exemplified by Mem|8. We propose a corrective framework grounded in a consistency condition requiring that the subconscious layer be a valid coarsening of raw trace equivalence; a stability theorem with explicit projection distortion; a bi-Lipschitz defect characterizing symmetric distortion from false merges and false splits; a bidirectional projection constraint preventing hallucination via ungrounded symbolic injection; a holonomy measure capturing path-dependent accumulated inconsistency; and a unified structural stability theorem showing that obstruction, distortion, and holonomy form an instability cascade. A variational principle unifies these as a constrained minimization problem, from which CLIO emerges as gradient flow. Operator realizations cover olfactory perception (receptor operators as topological classifiers, with spectral admissibility characterized via Katz-admissibility), bistable perception (obstruction-limited switching), Zeno dynamics (derived from first principles), and gamma-distributed dwell times (with a random-k mixture model recovering heavy-tailed behavior). An operator hierarchy makes the five-level stack from consistency axiom to observable statistics explicit. Sections on projection-first errors, admissible fragments, and media thermodynamics connect the apparatus to information ecosystems, biological life as a constraint-closure endomorphism, and pathology as sheaf failure. The architecture is reframed as timeful computation in the sense of Fant. The result is a framework in which the subconscious must exert less influence precisely when its internal coherence is least assured.

Contents

1 Introduction: From Suggestive Systems to Structural Architectures **6**

2 The Category Error: Representational Leakage and Hidden Priors **7**
2.1 Consequence 7

3 Minimal Architecture: Separation of Trace, Structure, and Influence **7**
3.1 Raw Trace Layer 8
3.2 Resonance Graph Layer 8
3.3 Meta-Wave Attractor Layer 8
3.4 Bias Layer (Subconscious) 8
3.5 Governing Constraint 8

4 RSVP Interpretation: Coarse-Graining Over Trajectories **8**
4.1 Mapping 9
4.2 Key Observation 9

5 The Failure Mode: Compression-Induced Obstruction **9**
5.1 Obstruction Score 10

6 The Consistency Condition **10**
6.1 Raw Trace Equivalence 10
6.2 Subconscious Equivalence 10
6.3 Main Condition 10
6.4 Failure Modes 10
6.5 The Frequency Mapping Failure 11

7 Stability Theorem with Projection Distortion **11**
7.1 Setup 11
7.2 Assumptions 11
7.3 Theorem 12
7.4 Interpretation 12

8 Obstruction–Distortion Relation **12**
8.1 Failure Cases 13
8.2 Connection to Stability 13

9 Bi-Lipschitz Projection Defect **13**
9.1 Refined Stability Condition 14

10 Bidirectional Projection Constraint **14**

11 Holonomy and Path-Dependent Obstruction **15**
11.1 Parallel Transport in Projection Space 15

11.2 Consequences 15
11.3 Connection to RSVP 16

12 Unified Structural Stability Theorem **16**
12.1 Setup 16

13 Variational Principle of Consistent Compression **17**
13.1 Energy Functional 17
13.2 Variational Problem 17
13.3 Regimes 17

14 Projection Before Coherence as a Universal Failure Mode **18**

15 Admissibility and the Thermodynamics of Fragmented Media **19**

16 Autonomy, Attention, and the Architecture Ordering Principle **19**
16.1 Architecture Ordering Principle 20

17 CLIO as Consistency Enforcement **20**
17.1 Definition 20
17.2 Role 20
17.3 Convergence and Limits 21

18 Implementation Discipline **21**
18.1 Stage 1: Trace Logging 21
18.2 Stage 2: Read-Only Summarization 21
18.3 Stage 3: Transparent Influence 21
18.4 Stage 4: Controlled Reinforcement 21
18.5 Principle 21

19 Olfaction as Energy-Selective Transition **22**
19.1 Vibrational Hypothesis 22
19.2 Isotope Sensitivity 22
19.3 Status and Working Assumption 22

20 RSVP Reinterpretation of Olfaction **23**
20.1 Odorant as Field Perturbation 23
20.2 Receptor as Constraint Surface 23
20.3 Resonance as Alignment 23
20.4 Activation as Entropy-Resolving Closure 23
20.5 Isotope Sensitivity Revisited 24

21 Olfaction as a Local Instance of Variational Consistency **24**
21.1 Local Functional 24
21.2 General Principle 24
21.3 Independence from Turin's Mechanism 25

22 Receptor Operators and Admissible Transition Classes **25**
22.1 Trajectory Space 25
22.2 Local Constraint Structure 25
22.3 Definition of the Receptor Operator 25
22.4 Connection to Variational Principle 26
22.5 Spectral Interpretation 26
22.6 Composition Across Receptors 26

23 Bistable Perception as Obstruction-Limited Dynamics **26**
23.1 RSVP Interpretation 27

24 Non-Commutativity and Temporal Nonlocality **27**
24.1 RSVP Correspondence 27

25 Zeno Dynamics and Obstruction Geometry **28**
25.1 Setup 28
25.2 Intrinsic Transition Time 28
25.3 Effect of Updates 28
25.4 Dwell Time 28
25.5 Connection to CLIO 29

26 Gamma-Distributed Dwell Times from Obstruction Dynamics **29**
26.1 Escape as a Multi-Step Process 29
26.2 Gamma Distribution 29
26.3 RSVP Interpretation of Parameters 30
26.4 Special Cases 30
26.5 Mean and Zeno Connection 30
26.6 CLIO and Irreducible Complexity 30
26.7 Cognitive Interpretation 30

27 Mixture Model: Random Obstruction Depth and Heavy Tails **30**
27.1 Random Depth Model 31
27.2 Log-Normal Emergence 31
27.3 RSVP Interpretation 31
27.4 CLIO and Irreducible Floor 31

28 Katz-Admissibility and Spectral Obstruction in Receptor Operators **32**
28.1 Katz-Admissibility 32
28.2 Connection to Obstruction Cohomology 32
28.3 Interpretation 33
28.4 Isotope Sensitivity as Ramification Change 33
28.5 Generalization 33

29 Operator Hierarchy: From Axiom to Realization **33**
29.1 The Five Levels 33
29.2 Vertical Connections 34

29.3 Summary 35

30 Timeful Computation and the Rejection of Primitive Sequentiality 35
30.1 Marine Transduction as Preservation of Dependency 35
30.2 RSVP as a Timeful Substrate 36
30.3 Rolling Coherence and Event Gating 36
30.4 CLIO as a Consistency Wavefront 37
30.5 TARTAN and Structural Concurrency 37
30.6 Conclusion 38

31 The Projection Error and the Clip Economy 38
31.1 The Projection Error 38
31.2 Optimization for Recoverability 39
31.3 The Clip Economy as Distributed Projection Error 39
31.4 Admissible Fragments and Reconstruction Pressure 40
31.5 Life as a Constraint-Closure Loop 40

32 Natural Implementation Patterns and Their Structural Consequences 41
32.1 Operationalizing Holonomy as a Frequency Threshold 42
32.2 Mapping Valence onto a Fixed Emotional Submanifold 42
32.3 Batch-Gating Recall Arbitration 43
32.4 The Common Structure 44

33 Conclusion 44

A Numerical Simulation of Obstruction and CLIO Repair 47
A.1 Trace Space 47
A.2 Obstruction Score 47
A.3 CLIO Repair Algorithm 48
A.4 Pseudocode 48
A.5 Interpretation 49

1. Introduction: From Suggestive Systems to Structural Architectures

This section situates the paper's central concern: what does it actually mean for a memory system to have a subconscious layer, and what goes wrong when that layer is built on conventions rather than derived from the system's own dynamics?

The recent emergence of recursive memory architectures has introduced systems that appear to exhibit subconscious-like behavior through internal summarization and pattern reinforcement. Among these, Mem|8 represents a compelling but incomplete attempt to construct a self-referential memory substrate.

However, the current formulation suffers from a critical ambiguity: it conflates evocative representational choices with derived structural properties. In particular, mappings such as frequency-encoded affect—associating emotional states with fixed numerical frequencies—introduce external coordinate systems that are not grounded in the internal dynamics of the system itself.

This paper argues that such mappings constitute a category error. They transform what should be an emergent structure into an imposed encoding, thereby compromising the generative integrity of the model. The critique is not stylistic: any fixed mapping from semantic states to numerical bases that is not derivable from the system's internal dynamics introduces a hidden prior. Once that is stated cleanly, the objection becomes structural rather than rhetorical.

We propose a minimal architectural correction grounded in constrained field dynamics and recursive coarse-graining, and demonstrate that once these constraints are enforced, the role of the subconscious becomes both sharply defined and formally tractable.

The argument proceeds as follows. Section 2 diagnoses the category error and its consequences. Section 3 introduces the minimal four-layer architecture required to avoid it. Section 4 interprets the system within the RSVP field-theoretic framework. Section 5 formalizes the failure mode of compression-induced obstruction. Section 6 states and proves the central consistency condition. Section 7 proves the stability theorem with explicit projection distortion. Section 8 establishes the Obstruction–Distortion relation linking the two. Section 17 defines the CLIO repair operator. Section 18 describes the correct implementation discipline. Sections 19–21 develop olfaction as a biological instantiation. Section 22 formalizes receptors as topological classifiers. Sections 23–26 extend to bistable perception. Section 27 introduces the random-k mixture model for heavy-tailed dwell times. Section 28 connects receptor spectral admissibility to Katz-admissibility and obstruction cohomology. Section 29 makes the five-level operator hierarchy explicit. Section 31 connects to the projection error, clip economy, and biological life. Section 30 covers timeful computation. Section 32 examines three natural implementation patterns and the structural consequences of each. Section 33 closes with irreducible obstruction.

2. The Category Error: Representational Leakage and Hidden Priors

Assigning fixed numbers to emotions—joy at 440 Hz, grief at 220 Hz—seems like a natural design choice. This section explains why it is not an implementation detail but a fundamental structural mistake, and what the correct requirement actually is.

A system that assigns fixed numerical representations to semantic states without derivation from its internal dynamics implicitly imports external structure. This introduces a hidden prior that cannot be audited or corrected from within the system.

Let $\mathcal{X}$ denote the space of raw experiential traces. A valid internal representation must arise from transformations

$$\pi : \mathcal{X} \to \mathcal{M}$$

that are functions of $\mathcal{X}$ alone.

Any mapping of the form

$$f : \text{semantic state} \to \mathbb{R}$$

that is externally specified—for example, assigning joy to 440 Hz or grief to 220 Hz—violates this condition.

2.1. Consequence

Such mappings do not compress or reveal structure; they overwrite it. The resulting system is no longer learning its own internal geometry but projecting onto a predefined coordinate system. The model stops being generative and becomes indexed to a culture, a convention, or an arbitrary initialization.

In formal terms: an externally specified frequency map defines equivalence classes in the subconscious space $\mathcal{S}$ that have no grounding in the recurrence structure of the raw trace space $\mathcal{M}$. These classes therefore have no valid preimage under the projection π. We formalize this as a violation of the Consistency Condition in Section 6.

3. Minimal Architecture: Separation of Trace, Structure, and Influence

If the category error identifies what not to do, this section identifies the minimal structure needed to do it correctly. Four layers are required, and the key constraint governing them is simple to state: the subconscious must whisper, not decide.

We propose that any valid recursive memory system must minimally decompose into four layers. The key constraint governing their interaction is stated at the end of this section.

3.1. Raw Trace Layer

The space $\mathcal{X}$ of recorded trajectories. This layer is strictly append-only and must preserve temporal and structural fidelity. It is the ground truth against which all higher-order structure is validated.

3.2. Resonance Graph Layer

A derived structure $G = (V, E)$ where nodes represent trace fragments and edges encode relational similarity or interaction, weighted by recency, retrieval co-occurrence, semantic distance, affective intensity, and successful reuse. This layer performs local organization without imposing global structure.

3.3. Meta-Wave Attractor Layer

A coarse-grained representation

$$\mathcal{A} : G \to \tilde{G}$$

that captures large-scale regularities and recurring motifs. This is the primary site of compression and the natural location of subconscious structure.

3.4. Bias Layer (Subconscious)

A low-amplitude modulation field

$$B : \tilde{G} \to \mathbb{R}$$

which influences retrieval and inference but does not determine outcomes.

3.5. Governing Constraint

The subconscious must whisper, not decide.

Formally, the influence of B must be bounded relative to the primary inference dynamics. This is not a heuristic but a stability condition, formalized in Section 7. When the bias layer becomes high-gain relative to the primary inference loop, it ceases to be a prior and becomes a controller, producing runaway attractors or hallucinated coherence.

4. RSVP Interpretation: Coarse-Graining Over Trajectories

The RSVP framework provides the field-theoretic language the paper works in. This section translates the memory architecture into that language and makes a point that is easy to miss: compression is not neutral. The meta-wave layer does not merely forget—it can actively create inconsistencies that were not there before.

Within the Relativistic Scalar-Vector Plenum (RSVP) framework, the system state is represented as a triple:

$$X(x,t) = (\Phi, \mathbf{v}, \mathcal{S})$$

where Φ encodes structural salience (constraint density), $\mathbf{v}$ encodes admissible transitions (trajectory flow), and $\mathcal{S}$ encodes configurational accessibility (entropy over futures).

Mem|8-like architectures can be interpreted as performing recursive coarse-graining over trajectory histories in this field.

4.1. Mapping

- Φ corresponds to structural salience (trace density and recurrence weight)
- $\mathbf{v}$ corresponds to transition dynamics between traces
- $\mathcal{S}$ corresponds to configurational accessibility (retrieval uncertainty)

The meta-wave layer is thus a projection

$$\Pi : X \to \tilde{X}$$

that reduces resolution while preserving large-scale structure. The subconscious layer is the residue-preserving gluing structure that helps future retrieval avoid total Markov collapse.

4.2. Key Observation

Compression is not neutral. While naive views assume that compression reduces information, in recursive systems it can introduce new inconsistencies not present in the original traces. The meta-wave layer does not merely forget; it can construct equivalences that did not previously exist. This is the site where CLIO enters as a repair operator: not optimization, but consistency repair over the induced structure.

5. The Failure Mode: Compression-Induced Obstruction

Here the central failure mode gets a name and a number. An obstruction is what happens when the subconscious layer either conflates things that should remain distinct or separates things that should be equivalent. Both directions are forms of hallucination.

Definition 5.1. *Let $\pi : \mathcal{M} \to \mathcal{S}$ be a coarse-graining operator on raw trace space. An* obstruction *arises when*

$$\pi(\gamma_i) = \pi(\gamma_j) \quad \textit{but} \quad \gamma_i \not\sim \gamma_j \textit{ in } \mathcal{M},$$

or when

$$\pi(\gamma_i) \neq \pi(\gamma_j) \quad \textit{but} \quad \gamma_i \sim \gamma_j \textit{ in } \mathcal{M}.$$

The first case is false coherence: distinct trajectories are merged without recurrence support. The second is false conflict: similar trajectories are separated. Both constitute hallucination.

5.1. Obstruction Score

Define the obstruction score:

$$\mathcal{O}(\pi) = \sum_{i,j} \left| \mathbf{1}_{\gamma_i \sim \gamma_j} - \mathbf{1}_{\pi(\gamma_i)=\pi(\gamma_j)} \right|.$$

A consistent projection satisfies $\mathcal{O}(\pi) = 0$.

6. The Consistency Condition

This is the paper's central theorem. The subconscious layer is valid if and only if it refines structure that already exists in the raw traces—it cannot invent its own equivalences from outside. The frequency-mapping failure from Section 2 is shown here to be a direct violation of this condition.

The validity of any subconscious or bias layer rests on a single structural requirement: it must not introduce distinctions or identifications that are not already supported by the raw trace dynamics.

6.1. Raw Trace Equivalence

Let $\mathcal{M}$ denote the space of raw traces with metric d, and let

$$R : \mathcal{M} \times \mathcal{M} \to \mathbb{R}$$

be a recurrence or similarity relation. This induces an equivalence relation $\sim$ on $\mathcal{M}$ via a threshold ε:

$$\gamma_i \sim \gamma_j \iff R(\gamma_i, \gamma_j) < \varepsilon.$$

6.2. Subconscious Equivalence

Let $\pi : \mathcal{M} \to \mathcal{S}$ be the projection into the subconscious space, and let $\sim_{\mathcal{S}}$ denote the induced equivalence:

$$\gamma_i \sim_{\mathcal{S}} \gamma_j \iff \pi(\gamma_i) = \pi(\gamma_j).$$

6.3. Main Condition

Theorem 6.1 (Consistency Condition)**.** *A subconscious layer $\mathcal{S}$ is consistent if and only if its induced equivalence relation is a coarsening of the raw trace equivalence:*

$$\forall \gamma_i, \gamma_j \in \mathcal{M}, \quad \gamma_i \sim_{\mathcal{S}} \gamma_j \;\Rightarrow\; \gamma_i \sim \gamma_j.$$

Equivalently, the fibers of π must be unions of equivalence classes of $\sim$.

6.4. Failure Modes

Two forms of violation arise.

False Coherence (Over-Merging). If $\gamma_i \not\sim \gamma_j$ but $\gamma_i \sim_{\mathcal{S}} \gamma_j$, the system introduces spurious similarity. Distinct trajectories are collapsed into a single class.

False Conflict (Over-Splitting). If $\gamma_i \sim \gamma_j$ but $\gamma_i \not\sim_{\mathcal{S}} \gamma_j$, the system introduces distinctions without support in the trace space.

Both constitute forms of hallucination: the subconscious layer is asserting structure that the raw experience does not support.

6.5. The Frequency Mapping Failure

Mappings such as

$$\text{“joy”} \mapsto 440\,\text{Hz}, \quad \text{“grief”} \mapsto 220\,\text{Hz}$$

define equivalence classes in $\mathcal{S}$ that have no grounding in the recurrence structure of $\mathcal{M}$. These mappings therefore violate Theorem 6.1 by introducing equivalence classes without valid preimages. This is the formal statement of the category error identified in Section 2.

7. Stability Theorem with Projection Distortion

The consistency condition has a dynamical consequence: bad compression does not just produce wrong answers, it shrinks the safe operating region of the entire system. This section makes that precise. The projection is not a neutral operation—its distortion directly limits how much the subconscious layer is allowed to influence behavior.

7.1. Setup

Let $\mathcal{M}$ be the raw trace space and $\pi : \mathcal{M} \to \mathcal{S}$ the subconscious projection. We model iterated retrieval dynamics as:

$$x_{t+1} = F(x_t) + \lambda B(\pi(x_t)),$$

where F is the primary inference map, B is the bias operator, and $\lambda \geq 0$ controls subconscious influence.

Irreducible obstruction introduces a nonzero lower bound on L_π, thereby imposing an upper bound on admissible gain λ even under optimal repair.

7.2. Assumptions

(A1) Contractive Primary Dynamics.

$$\|F(x) - F(y)\| \leq L_F \|x - y\|, \quad 0 \leq L_F < 1.$$

(A2) Projection Distortion Bound. The projection π is Lipschitz with constant L_π:

$$\|\pi(x) - \pi(y)\| \leq L_\pi \|x - y\|.$$

This condition is satisfied if π is a quotient map with bounded fiber diameter. Violation of the Consistency Condition (Theorem 6.1) increases L_π or renders it undefined.

(A3) Bias Operator Regularity.

$$\|B(u) - B(v)\| \leq L_B \|u - v\|, \quad u, v \in \mathcal{S}.$$

7.3. Theorem

Theorem 7.1 (Stability)**.** *The iterated memory dynamics are stable if*

$$L_F + \lambda L_B L_\pi < 1.$$

Proof. Let x_t and y_t be two admissible trajectories. Then:

$$\begin{aligned}
\|x_{t+1} - y_{t+1}\| &= \|F(x_t) - F(y_t) + \lambda(B(\pi(x_t)) - B(\pi(y_t)))\| \\
&\leq \|F(x_t) - F(y_t)\| + \lambda \|B(\pi(x_t)) - B(\pi(y_t))\| \\
&\leq L_F \|x_t - y_t\| + \lambda L_B \|\pi(x_t) - \pi(y_t)\| \\
&\leq (L_F + \lambda L_B L_\pi) \, \|x_t - y_t\|.
\end{aligned}$$

Thus the system is contractive when $L_F + \lambda L_B L_\pi < 1$. □

7.4. Interpretation

The projection π is not neutral. If it distorts geometry (large L_π), then even small subconscious gain λ can destabilize the system. Violation of the Consistency Condition manifests dynamically as an increase in L_π, tightening the allowable bound on subconscious influence.

Corollary 7.2 (Geometric Hallucination Threshold)**.** *If*

$$L_F + \lambda L_B L_\pi \geq 1,$$

then subconscious influence amplifies perturbations. Instability arises not merely from strong bias, but from distortion introduced by invalid coarse-graining. Bad compression does not merely create wrong answers—it shrinks the stability region of the entire system.

8. Obstruction–Distortion Relation

This section connects the counting measure of inconsistency (obstruction score) to the geometric measure of distortion (Lipschitz constant). The bridge between them is what allows the appendix simulation results to feed directly into the stability theorem.

We now relate the obstruction score $\mathcal{O}(\pi)$ to the geometric distortion L_π, completing the bridge between the appendix simulation and the stability theorem.

Proposition 8.1 (Obstruction–Distortion)**.** *There exists a monotone function f such that*

$$L_\pi \geq f\left(\frac{\mathcal{O}(\pi)}{|\mathcal{M}|^2}\right).$$

Sketch. Each violating pair (γ_i, γ_j) contributes a local inconsistency between the metric structure of $\mathcal{M}$ and the partition induced by π. To realize such inconsistencies, π must either collapse large distances (false merge) or expand small distances (false split). In either case, the supremum defining L_π increases. Since $\mathcal{O}(\pi)$ counts violating pairs, normalized obstruction provides a lower bound on the minimal distortion required. □

8.1. Failure Cases

False Merges. $d(\gamma_i, \gamma_j) \gg \varepsilon$ but $\pi(\gamma_i) = \pi(\gamma_j)$ forces large compression of metric structure.

False Splits. $d(\gamma_i, \gamma_j) < \varepsilon$ but $\pi(\gamma_i) \neq \pi(\gamma_j)$ forces expansion of local neighborhoods.

Both effects increase L_π.

Corollary 8.2. *If $\mathcal{O}(\pi) = 0$, then π can be chosen to be non-expansive: $L_\pi \leq 1$. Consistency implies geometric stability of the projection.*

8.2. Connection to Stability

Substituting Proposition 8.1 into Theorem 7.1:

$$L_F + \lambda L_B f\left(\frac{\mathcal{O}(\pi)}{|\mathcal{M}|^2}\right) < 1.$$

Key Consequence. Obstruction directly reduces the allowable subconscious gain λ. As $\mathcal{O}(\pi)$ increases, the system becomes more sensitive to subconscious influence and more prone to instability.

Remark 8.3. *Subconscious hallucination is not merely representational error; it is a geometric deformation of the trace space that propagates into dynamical instability.*

9. Bi-Lipschitz Projection Defect

The stability analysis so far only captures one direction of distortion—expansion. But compression can also collapse distances that should remain distinct. This section introduces a symmetric measure that captures both, and shows they are dual faces of the same geometric defect.

The stability analysis of Section 7 bounds projection distortion via a Lipschitz constant L_π, capturing expansion of distances. However, obstruction arises from both expansion (false splitting) and collapse (false merging). A symmetric treatment requires a bi-Lipschitz characterization.

Definition 9.1 (Bi-Lipschitz Defect)**.** *Let $\pi : \mathcal{M} \to \mathcal{S}$. Define the bi-Lipschitz defect δ_π as the smallest $\delta \geq 0$ such that for all $x, y \in \mathcal{M}$:*

$$(1 - \delta)\, d(x, y) \leq d(\pi(x), \pi(y)) \leq (1 + \delta)\, d(x, y).$$

$\delta_\pi = 0$ corresponds to an isometric projection (perfect consistency). Positive δ_π captures both contraction and expansion symmetrically.

Proposition 9.2 (Symmetric Obstruction–Distortion)**.** *There exists a monotone function g such that:*

$$\delta_\pi \geq g\left(\frac{\mathcal{O}(\pi)}{|\mathcal{M}|^2}\right).$$

9.1. Refined Stability Condition

The stability theorem becomes:

$$L_F + \lambda L_B(1 + \delta_\pi) < 1.$$

False merges (collapse) and false splits (expansion) are not distinct pathologies but dual manifestations of the same geometric defect. The system's instability is governed by total distortion, not just expansion.

Remark 9.3. *The bi-Lipschitz defect δ_π provides a unified geometric measure of subconscious hallucination.*

10. Bidirectional Projection Constraint

The consistency condition governs compression going downward—from traces into the subconscious. But symbols can also be injected upward into the field from outside. Without constraints on that direction, a system can hallucinate not by compressing badly but by believing things it could never have perceived.

Projection from the RSVP field into symbolic representation is well-defined:

$$\Pi : X \to \Sigma.$$

However, symbolic inputs must also be reintroduced into the field:

$$\iota : \Sigma \to X.$$

Without constraint on ι, symbolic inputs can inject structure into X that has no grounding in the trace space $\mathcal{M}$, bypassing the Consistency Condition entirely.

Definition 10.1 (Ungrounded Injection)**.** *An injection ι is ungrounded if there exists $\sigma \in \Sigma$ such that $\iota(\sigma) \notin \mathrm{Im}(\Pi \circ \mathcal{M})$, i.e., it has no valid preimage in the sensory-derived trace space.*

Theorem 10.2 (Bidirectional Consistency)**.** *A projection system is fully consistent if: (1)* Π *satisfies Theorem 6.1, and (2)* $\iota(\Sigma) \subseteq \mathrm{Im}(\Pi \circ \mathcal{M})$.

Symbols may only re-enter the field if they correspond to structures that could have been generated by the system's own perceptual pipeline. Violation produces two hallucination classes:

- **Forward hallucination:** invalid compression (Sections 5–8)
- **Backward hallucination:** invalid injection (this section)

A system hallucinates not only when it compresses incorrectly, but when it believes what it could never have perceived.

11. Holonomy and Path-Dependent Obstruction

Local consistency is not enough. A system can be locally consistent at every point while accumulating structural damage along trajectories—the cognitive equivalent of a path that returns to where it started but arrives wrong. Holonomy is the measure of that accumulated error.

Local consistency does not guarantee global consistency. Obstruction may accumulate along trajectories even when all local equivalence conditions are satisfied.

11.1. Parallel Transport in Projection Space

Consider transporting an equivalence class along a loop γ in trace space:

$$\gamma : [0, 1] \to \mathcal{M}, \quad \gamma(0) = \gamma(1).$$

Let π_t denote the evolving projection along γ.
Definition 11.1 (Holonomy Defect)**.** *The holonomy defect along* γ *is:*

$$H(\gamma) = d\big(\pi_1(\gamma(0)),\, \pi_0(\gamma(0))\big).$$

$H(\gamma) = 0$ indicates global consistency; $H(\gamma) > 0$ indicates path-dependent inconsistency. Crucially, holonomy defect can be nonzero even when $\mathcal{O}(\pi)$ is locally minimized.

11.2. Consequences

Holonomy captures failures that obstruction alone cannot:

- *Obstruction*: pairwise inconsistency between traces,
- *Holonomy*: accumulated inconsistency along trajectories.

A monitoring operator $\mathcal{H} : \gamma \mapsto H(\gamma)$ detects structural damage invisible to local repair.

Curvature measures local error; holonomy measures memory of error.

11.3. Connection to RSVP

Holonomy corresponds to nontrivial circulation in the vector field $\mathbf{v}$:

$$\oint_\gamma \mathbf{v} \cdot d\mathbf{x} \neq 0.$$

Path-dependent cognitive inconsistency is directly measurable as field circulation. This connects the Ambrose-Singer Custodian (which measures local curvature against global holonomy) to the RSVP vector field as its natural physical substrate.

12. Unified Structural Stability Theorem

The three failure modes—pairwise inconsistency, geometric distortion, and path-dependent drift—are not independent problems. They form a cascade: each one forces the next. This section states that chain as a single theorem and makes explicit what it means for a system to be globally stable.

We unify the three failure modes of recursive memory compression — pairwise inconsistency (obstruction), geometric distortion (bi-Lipschitz defect), and path-dependent inconsistency (holonomy) — into a single structural constraint governing dynamical stability.

12.1. Setup

Let $\hat{\mathcal{O}} = \mathcal{O}(\pi)/|\mathcal{M}|^2$ denote normalized obstruction. There exist monotone functions f, g such that:

$$\delta_\pi \geq f(\hat{\mathcal{O}}), \tag{1}$$

$$\sup_\gamma H(\gamma) \geq g(\delta_\pi). \tag{2}$$

Thus $\sup_\gamma H(\gamma) \geq (g \circ f)(\hat{\mathcal{O}})$: local inconsistency forces geometric distortion; geometric distortion accumulates into path-dependent inconsistency.

Theorem 12.1 (Unified Structural Stability)**.** *The system is globally stable only if:*

$$L_F + \lambda L_B(1 + \delta_\pi) < 1,$$

with $\delta_\pi \geq f(\hat{\mathcal{O}})$ *and* $\sup_\gamma H(\gamma) \geq g(\delta_\pi)$.

Corollary 12.2 (Instability Cascade)**.** *If* $\hat{\mathcal{O}} > 0$*, then:*

$$\hat{\mathcal{O}} > 0 \;\Rightarrow\; \delta_\pi > 0 \;\Rightarrow\; \sup_\gamma H(\gamma) > 0 \;\Rightarrow\; \textit{reduced admissible } \lambda.$$

There is no isolated failure mode. Structural inconsistency propagates across all levels.

In the irreducible case, $\hat{\mathcal{O}} \geq \hat{\mathcal{O}}^* > 0$ yields $\delta_\pi \geq \delta_{\min}$ and:

$$\lambda < \frac{1 - L_F}{L_B(1 + \delta_{\min})}.$$

Remark 12.3. *Subconscious hallucination is not a local defect but a structural cascade: it begins as an invalid equivalence, becomes a geometric deformation, accumulates as path-dependent inconsistency, and manifests as dynamical instability. Consistency is not a local property of representation but a global invariant. When it fails, it fails everywhere.*

13. Variational Principle of Consistent Compression

All of the conditions developed so far can be unified into a single energy functional that the system is trying to minimize. CLIO is not a procedure imposed from outside—it is the gradient flow of that functional. The subconscious emerges as the solution to a constrained variational problem, not as a design choice.

The preceding analysis can be unified by interpreting subconscious compression as a constrained variational problem over projection maps.

13.1. Energy Functional

Each projection $\pi : \mathcal{M} \to \mathcal{S}$ induces an obstruction score $\mathcal{O}(\pi)$, a geometric distortion δ_π, and a holonomy defect functional $H(\pi)$. Define the total compression energy:

$$\mathcal{E}(\pi) = \mathcal{O}(\pi) + \alpha\, \delta_\pi^2 + \beta\, \mathbb{E}_\gamma[H(\gamma)^2],$$

where $\alpha, \beta > 0$ weight geometric and path-dependent penalties, and the expectation is taken over a representative ensemble of loops γ in $\mathcal{M}$.

13.2. Variational Problem

Definition 13.1 (Consistent Compression)**.** *A projection π^* is admissible if it minimizes $\mathcal{E}(\pi)$ subject to the stability constraint:*

$$\pi^* = \arg\min_\pi \mathcal{E}(\pi) \quad \textit{such that} \quad L_F + \lambda L_B(1 + \delta_\pi) < 1.$$

Introducing a Lagrange multiplier $\mu \geq 0$:

$$\mathcal{L}(\pi) = \mathcal{E}(\pi) + \mu\left(L_F + \lambda L_B(1 + \delta_\pi) - 1\right),$$

critical points satisfy $\nabla_\pi \mathcal{L} = 0$. CLIO emerges as gradient flow on this functional:

$$\partial_t \pi = -\nabla_\pi \mathcal{L}.$$

Repair is therefore not procedural but variational: the system continuously descends toward consistency under stability constraints.

13.3. Regimes

Fully Consistent Regime. If $\mathcal{O}(\pi) = 0$, $\delta_\pi = 0$, $H(\gamma) = 0$, then π^* is an isometric quotient and stability is maximized.

Irreducible Regime. If $\inf_\pi \mathcal{O}(\pi) = \mathcal{O}^* > 0$, then $\delta_\pi \geq \delta_{\min} > 0$ and:

$$\lambda < \frac{1 - L_F}{L_B(1 + \delta_{\min})}.$$

The system cannot eliminate inconsistency; it can only trade off distortion against stability.

Dual Interpretation. $\mathcal{E}$ admits three readings: $\mathcal{O}(\pi)$ measures information inconsistency, δ_π measures geometric strain, and $H(\gamma)$ measures historical inconsistency. Minimizing $\mathcal{E}$ enforces alignment across all three simultaneously.

Theorem 13.2 (Variational Stability Principle)**.** *A recursive cognitive system is stable if and only if its subconscious projection π is a constrained minimizer of $\mathcal{E}$. Failure to minimize $\mathcal{E}$ produces a cascade:*

$$\textit{obstruction} \to \textit{distortion} \to \textit{holonomy} \to \textit{instability}.$$

Remark 13.3. *Subconscious processing is not an auxiliary mechanism but the solution to a constrained variational problem over representations.*

14. Projection Before Coherence as a Universal Failure Mode

The tendency to compress before establishing coherence turns out to be a universal pattern, appearing in signal processing, machine learning, and media systems alike. This section names it as a single failure mode and shows that it always introduces obstruction, regardless of domain.

A unifying failure mode across sensory processing, machine learning, and sociotechnical media systems is the premature application of projection prior to structural coherence.

Definition 14.1 (Projection-First System)**.** *A system is projection-first if there exists a neighborhood $U \subset \mathcal{M}$ such that $\pi|_U$ is applied before verifying that the recurrence-induced equivalence relation $\sim$ is locally consistent on U.*

Proposition 14.2 (Universality of Projection Error)**.** *Any projection-first system necessarily introduces nonzero obstruction $\mathcal{O}(\pi) > 0$ unless the underlying trace space is trivial.*

This failure manifests across domains: signal processing (PCM discretization projects continuous waveforms before preserving micro-temporal structure), machine learning (vector embeddings collapse trajectories into static coordinates), and media systems (fragments are extracted without preserving compositional structure).

Coherence must be established prior to projection.

15. Admissibility and the Thermodynamics of Fragmented Media

Not all fragments are equal. An admissible fragment is one that carries enough of its original context to be reintegrated without fabricating structure. This section develops a thermodynamic account of why high-throughput media systems tend to select for the wrong kind of fragment, and what an admissibility constraint would actually require.

We model information ecosystems as driven dissipative systems over a space of fragments $\mathcal{F}$. Each fragment $f \in \mathcal{F}$ is a projection of some underlying trace $\gamma \in \mathcal{M}$.

Definition 15.1 (Admissible Fragment)**.** *A fragment f is admissible if: (1) $\pi_f(\gamma)$ is locally coherent (extraction validity); (2) for compatible fragments f_1, f_2, their joint reconstruction increases information: $I(f_1 \cup f_2) > \max(I(f_1), I(f_2))$ (compositional gain); (3) reconstruction is stable under small permutations (order sensitivity).*

We define entropy S as contextual loss under projection. In high-throughput environments, fragments with higher entropy propagate more efficiently:

$$\frac{d\phi}{dt} \propto -\frac{\partial S}{\partial f},$$

where ϕ denotes engagement potential.

Proposition 15.2 (Entropy Selection Principle)**.** *In a throughput-optimized system, fragments with higher entropy S are preferentially selected.*

This induces an *ontological inversion*: fragments cease to refer to global structure and become terminal objects. Degenerate fragments arise as violations of admissibility: terminal fragments (no reconstruction path exists), drift fragments (projection maps change across propagation), and adversarial fragments (no valid preimage exists in $\mathcal{M}$).

> *Modern media systems can be characterized as entropy-maximizing projection systems operating without admissibility constraints.*

16. Autonomy, Attention, and the Architecture Ordering Principle

Autonomy—the capacity to rely on one's own compressed history—is not a binary property but a quantity bounded by the quality of the system's own consistency. The more inconsistency remains unresolved, the less the system can safely trust itself. This section also states the fundamental ordering: preserve, then enforce, then repair, then project.

From the stability theorem, the weight of internally generated influence λ is bounded by:

$$\lambda < \lambda_{\max} = \frac{1 - L_F}{L_B f(\mathcal{O}^*)}.$$

Autonomy is inversely bounded by unresolved inconsistency.

In cognitive terms, this formalizes the necessity of reducing reliance on internally generated interpretations when underlying experiences are contradictory. Failure to satisfy this constraint leads to runaway feedback loops and pathological attractors.

16.1. Architecture Ordering Principle

The structural requirements of a coherent system form an ordering over operators:

$$\text{Preserve} \rightarrow \text{Enforce} \rightarrow \text{Repair} \rightarrow \text{Project}$$

1. **Preserve:** maintain raw dependency relations (Marine Algorithm),
2. **Enforce:** ensure local consistency (TARTAN sheaf gluing),
3. **Repair:** resolve global obstructions (CLIO),
4. **Project:** map to reduced representation (interface layer).

Proposition 16.1. *Any reordering that applies projection prior to repair introduces irreversible information loss.*

Projection is the final step, not the first.

17. CLIO as Consistency Enforcement

CLIO is the repair operator: the mechanism that detects and corrects violations of the consistency condition. This section defines it precisely, explains what it can and cannot do, and introduces the concept of irreducible obstruction—the formal signature of genuine contradiction in experience.

To maintain consistency, a repair operator is required. The Constraint-Leveraged Inference Operator (CLIO) serves this role.

17.1. Definition

CLIO acts as:

$$\mathcal{C} : \pi \mapsto \pi'$$

minimizing obstruction across the quotient structure while preserving valid equivalence classes.

17.2. Role

1. Detect inconsistencies introduced by compression
2. Resolve conflicts via minimal structural adjustment
3. Preserve global coherence without reverting to raw trace space

CLIO is therefore not an optimization mechanism but a consistency-preserving operator. Its necessity follows from the structure of the system: under recursive application of π, obstructions accumulate unless actively repaired.

17.3. Convergence and Limits

CLIO is a greedy descent on $\mathcal{O}(\pi)$ and is not guaranteed to reach a global minimum. It may converge to a local minimum. This is not a failure of the operator but a structural feature of the trace space.

Irreducible Obstruction. If $\mathcal{O}(\pi^*) > 0$ at termination, this may reflect *irreducible obstruction*: the raw trace space itself contains incompatible recurrence structure that cannot be jointly quotiented without contradiction.

Irreducible obstruction corresponds to cognitive dissonance: the system has experienced trajectories that cannot be made mutually consistent by any admissible compression. Residual $\mathcal{O} > 0$ is therefore not an error state but a structurally meaningful condition. It is the formal signature of genuine contradiction in experience.

18. Implementation Discipline

The theory has a direct practical consequence: the system must be built in a specific order, with each layer validated before the next is activated. Skipping this sequencing is not a shortcut—it is the mechanism by which bias layers become trusted before their consistency has been established.

A correct system must be built in stages, with each layer validated before the next is made operative.

18.1. Stage 1: Trace Logging

Raw traces are recorded without transformation. Append-only. No inference.

18.2. Stage 2: Read-Only Summarization

The meta-layer is constructed but not used for decision-making. Summaries are observable but inert.

18.3. Stage 3: Transparent Influence

The bias layer is introduced with measurable, bounded effect. The influence of B is reported alongside retrieval outputs.

18.4. Stage 4: Controlled Reinforcement

Feedback is allowed only after validation of earlier stages.

18.5. Principle

Each layer must be falsifiable before being made operative. This sequencing prevents the failure mode where a bias layer is trusted before its consistency has been established.

19. Olfaction as Energy-Selective Transition

Turin's vibrational hypothesis of olfaction—that receptors detect molecular vibration frequencies, not just shape—provides the paper's key biological case study. Whether or not the hypothesis is empirically correct, it articulates a model of energy-selective activation that the RSVP framework can formalize rigorously. This section presents that model and its experimental status.

The dominant biological model of olfaction is shape-based recognition via receptor binding. An alternative proposal, due to Luca Turin, suggests that molecular vibrations may play a role in receptor activation alongside geometric matching.

19.1. Vibrational Hypothesis

The core physical relation is:

$$E = h\nu,$$

linking vibrational frequency ν to energy E. Turin's proposal is that receptor activation involves *inelastic electron tunneling*: an electron can tunnel across a receptor site only if the energy gap matches a vibrational mode of the bound molecule.

Under this model, the odorant does not merely bind geometrically. It enables a transition condition. The receptor measures energy compatibility in addition to shape, performing something analogous to local spectroscopy—sensitive to energy loss during electron transfer rather than optical absorption.

19.2. Isotope Sensitivity

Isotope substitution preserves molecular geometry while altering atomic mass, which changes vibrational frequencies without changing shape. If vibrational matching is operative, two isotopically distinct molecules with identical shapes could activate receptors differently. Some experimental evidence suggests isotope-dependent smell; other attempts at replication have been inconclusive.

19.3. Status and Working Assumption

The vibrational hypothesis remains controversial. The mainstream view retains shape-based G-protein-coupled receptor mechanisms as primary. Whether vibrational coupling contributes meaningfully is an open question.

Working Assumption. Regardless of its empirical status, the vibrational hypothesis provides a minimal model of *energy-selective activation*: a system that gates a signal not on shape alone but on the compatibility of an incoming perturbation with a local energy transition. This structure is sufficient for the theoretical development that follows, and does not depend on Turin's mechanism being the dominant pathway.

20. RSVP Reinterpretation of Olfaction

Rather than asking whether a molecule matches a receptor's shape or frequency, the RSVP framework asks whether the molecule's trajectory can be consistently integrated into the receptor's constraint structure. Activation is closure, not detection. This reframing makes isotope sensitivity a natural consequence rather than a brute empirical fact.

We reinterpret olfactory activation as a local constrained transition in the RSVP field $X = (\Phi, \mathbf{v}, S)$.

20.1. Odorant as Field Perturbation

An odorant molecule induces a structured perturbation in the field:

- Φ: local potential shaped by electronic structure
- $\mathbf{v}$: trajectory of approach, binding, and interaction
- S: accessibility of admissible transitions

The odorant is therefore not an object but a trajectory in field space carrying internal spectral structure.

20.2. Receptor as Constraint Surface

Each receptor defines a constraint manifold $\Sigma \subset X$. Admissible configurations lie on Σ; incompatible perturbations generate high obstruction and fail to complete a transition. Binding is an attempt to project an incoming trajectory onto Σ.

20.3. Resonance as Alignment

Turin's vibrational matching condition corresponds to a coherence requirement between flow and structure:

$$\langle \mathbf{v}, \nabla\Phi \rangle \;\approx\; \text{admissible transition energy.}$$

Only trajectories aligned with the local gradient structure of Φ can dissipate energy along permitted directions and complete the transition. This replaces "frequency detection" with alignment in field space: the system is not measuring a property of the molecule but testing whether the trajectory satisfies a local coherence condition.

20.4. Activation as Entropy-Resolving Closure

Receptor activation occurs when an incoming perturbation resolves a local mismatch:

1. the perturbation enters with high local obstruction,
2. an admissible transition pathway exists,
3. entropy decreases locally as the constraint is satisfied,
4. a stable transition state is reached.

Perception is the existence of a low-obstruction closure under constraint.

This is precisely the CLIO dynamic at the receptor scale: a perturbation either fails (no closure, no signal) or resolves (closure, perception event).

20.5. Isotope Sensitivity Revisited

Isotope substitution in this framework:

- leaves Σ unchanged (same geometry, same constraint surface),
- alters the spectral structure of the incoming trajectory (different vibrational modes, different admissible transitions).

Thus $\langle \mathbf{v}_{\text{isotope}}, \nabla\Phi \rangle$ may fail the coherence condition where $\langle \mathbf{v}, \nabla\Phi \rangle$ succeeds. The system is not detecting mass or shape; it is detecting whether a trajectory can complete a low-entropy transition. Isotope sensitivity follows as a consequence of that sensitivity, not as a brute fact about receptors.

21. Olfaction as a Local Instance of Variational Consistency

Olfaction is not a special case—it is a biological instance of the same variational principle that governs memory compression, concept recognition, and token acceptance. A receptor does not detect a molecule; it tests whether a molecular trajectory solves a local admissibility problem. This section makes that identification explicit.

The olfactory process provides a concrete biological instantiation of the variational principle underlying the broader framework.

21.1. Local Functional

At a receptor, define a local functional over incoming trajectories γ:

$$\mathcal{E}_{\text{rec}}(\gamma) = \mathcal{O}_{\text{local}}(\gamma) + \alpha\, \delta_{\text{local}}(\gamma)^2,$$

where $\mathcal{O}_{\text{local}}$ measures local obstruction and δ_{local} measures deviation from the constraint surface Σ.

Activation occurs if:

$$\exists\, \gamma \text{ such that } \mathcal{E}_{\text{rec}}(\gamma) \approx 0.$$

The receptor is not detecting a property of the molecule; it is solving a local minimization problem: does a low-obstruction path exist that integrates the incoming perturbation into the constraint structure?

21.2. General Principle

This example instantiates the detection principle stated throughout this paper:

> *Systems do not detect objects. They detect whether perturbations admit low-energy integration into their constraint structure.*

The same principle governs sensory perception (olfaction), cognition (concept recognition), symbolic systems (token acceptance under embedding constraints), and physical interaction (energy transitions in quantum systems). Olfaction is not exceptional but exemplary.

21.3. Independence from Turin's Mechanism

The variational structure extracted here does not depend on electron tunneling being the biological mechanism. Whether shape, vibration, or some combination governs actual olfactory transduction, the formal claim is the same: what the receptor is doing, at the level of information processing, is testing trajectory admissibility against a constraint surface. Turin's hypothesis is one physical realization of that structure. It need not be the only one.

22. Receptor Operators and Admissible Transition Classes

This section gives the receptor a precise mathematical form: it is a topological classifier that tests whether an incoming perturbation belongs to a homotopy class with a low-obstruction representative. A smell is not a property of a molecule; it is a pattern of satisfied constraints across a family of such operators.

We formalize the receptor as a local operator acting on trajectories in the RSVP field, completing the transition from analogy to native instance.

22.1. Trajectory Space

Let Γ denote the space of admissible trajectories:

$$\gamma : [0,1] \to X = (\Phi, \mathbf{v}, S).$$

Each incoming perturbation defines a trajectory $\gamma \in \Gamma$ localized near a receptor.

22.2. Local Constraint Structure

A receptor defines a constraint region $\Sigma \subset X$ with associated obstruction functional $\mathcal{O}_\Sigma(\gamma)$. Low values correspond to trajectories compatible with the receptor's structure.

22.3. Definition of the Receptor Operator

Definition 22.1 (Receptor Operator)**.**

$$\mathcal{R}_\Sigma : \Gamma \to \{0,1\}$$

is defined by:

$$\mathcal{R}_\Sigma(\gamma) = \begin{cases} 1 & \text{if } \exists\, \tilde{\gamma} \sim \gamma \text{ such that } \mathcal{O}_\Sigma(\tilde{\gamma}) \leq \varepsilon, \\ 0 & \text{otherwise,} \end{cases}$$

where $\tilde{\gamma} \sim \gamma$ denotes homotopy equivalence within the local region.

Activation does not depend on the exact trajectory, but on its homotopy class: if a low-obstruction representative exists, the receptor activates; otherwise the perturbation is rejected. Receptors are therefore topological classifiers, not geometric detectors.

22.4. Connection to Variational Principle

The operator $\mathcal{R}_\Sigma$ implements a local feasibility check for the variational problem:

$$\exists\, \gamma \quad \text{such that} \quad \mathcal{E}_{\text{rec}}(\gamma) \approx 0.$$

It is a binary projection of the underlying minimization process.

22.5. Spectral Interpretation

In the vibrational hypothesis, admissibility depends on energy matching: $\Delta E \in \text{Spec}(\Sigma)$. In RSVP terms, this is a condition on allowable transitions in Φ along γ. Spectral sensitivity is therefore a special case of admissible trajectory structure.

22.6. Composition Across Receptors

Let $\{\Sigma_i\}$ be a family of receptors. The olfactory system computes:

$$\mathcal{R}(\gamma) = \big(\mathcal{R}_{\Sigma_1}(\gamma), \ldots, \mathcal{R}_{\Sigma_n}(\gamma)\big).$$

Perception corresponds to the pattern of activated constraint operators.

> *A smell is not a property of a molecule, but a pattern of satisfied constraints.*

Proposition 22.2. *Receptor activation is equivalent to the existence of a homotopy class of trajectories admitting a low-obstruction representative under local constraints.*

Remark 22.3. *The system does not measure properties of inputs. It tests whether inputs can be consistently integrated into its constraint structure.*

Synthesis. Across all domains, a system perceives not what is present, but what can be made consistent.

23. Bistable Perception as Obstruction-Limited Dynamics

The Necker cube flips not because a decision is made, but because the system traverses a high-obstruction region between two nearly equivalent stable states. Perceptual switching is a trajectory, not a jump. This section develops that picture and connects switching dynamics to the obstruction gradient.

Bistable perception provides a concrete empirical instance of constraint-driven trajectory dynamics that connects directly to the obstruction framework.

In ambiguous figures such as the Necker cube, the perceptual system alternates between two stable interpretations without any change in external input. Experimental data indicate that transitions are preceded by continuous destabilization phases and exhibit nontrivial temporal structure. Quantum-inspired harmonic oscillator models demonstrate that perceptual states evolve continuously and only appear discrete upon observation, with transitions mediated by traversal of a potential barrier rather than discrete jumping.

This suggests that perceptual switching is not a jump between states but a traversal of an intermediate high-obstruction region in state space.

23.1. RSVP Interpretation

Within the RSVP field $X = (\Phi, \mathbf{v}, S)$:

- stable percepts correspond to local minima of $\mathcal{O}$,
- ambiguity corresponds to near-degenerate minima,
- switching corresponds to trajectories crossing a barrier in Φ.

Bistable perception is thus a dynamical process governed by obstruction gradients rather than discrete state selection.

24. Non-Commutativity and Temporal Nonlocality

The order in which perceptual updates occur changes the outcome—observation A followed by B is not the same as B followed by A. This section formalizes that path-dependence and connects it to the RSVP field structure, where the direction of the vector field $\mathbf{v}$ records the history of the trajectory.

A key feature of bistable perception is that successive observations do not commute: the order of perceptual updates influences the resulting state. Formally, for operations A and B acting on a perceptual state q:

$$A(B(q)) \neq B(A(q)).$$

This non-commutativity introduces temporal nonlocality: the system's state depends on its history of observations, not merely its current configuration.

24.1. RSVP Correspondence

Non-commutativity arises in the RSVP framework when trajectory-dependent updates alter the admissible region of X:

- $\mathbf{v}$ encodes trajectory direction and order dependence,
- Φ encodes the constraint landscape that updates modify,
- S encodes accessible future trajectories after each update.

Perception is therefore path-dependent:

The system does not occupy a state; it occupies a history-constrained trajectory.

25. Zeno Dynamics and Obstruction Geometry

The quantum Zeno effect—frequent observation stabilizes a state by preventing transitions—has an analogue in obstruction dynamics that requires no quantum mechanics to derive. This section shows that the characteristic dwell-time relation follows directly from the geometry of the obstruction barrier and the rate of coherence enforcement.

We derive the characteristic dwell-time relation observed in bistable perception from obstruction-limited dynamics in the RSVP framework, without invoking quantum postulates.

25.1. Setup

Consider two near-degenerate minima of $\mathcal{O}$ corresponding to perceptual states A and B, separated by a barrier of height $\Delta\mathcal{O}$. The system evolves under two competing processes: continuous drift under the gradient flow

$$\partial_t \pi = -\nabla_\pi \mathcal{O}(\pi),$$

and discrete coherence-enforcing updates occurring at interval ΔT.

25.2. Intrinsic Transition Time

In the absence of updates, escape time scales as:

$$t_0 \sim \frac{1}{\|\nabla_\pi \mathcal{O}\|} \exp\left(\frac{\Delta\mathcal{O}}{\sigma}\right),$$

where σ represents effective fluctuation strength induced by entropy S.

25.3. Effect of Updates

Each update enforces partial projection consistency, resetting the system toward a local minimum. For small ΔT, escape probability per interval is:

$$p \sim \left(\frac{\Delta T}{t_0}\right)^2,$$

which is quadratic because escape requires both sufficient drift along $\nabla_\pi \mathcal{O}$ and persistence across the update boundary.

25.4. Dwell Time

The mean dwell time then satisfies:

$$\langle T \rangle = \langle N \rangle \Delta T \sim \frac{t_0^2}{\Delta T}.$$

$$\boxed{\langle T \rangle = \frac{t_0^2}{\Delta T}}$$

This reproduces the Zeno-like relation observed empirically without quantum mechanics. The three terms carry direct RSVP interpretations: t_0 is the intrinsic obstruction-limited transition time, ΔT is the coherence enforcement interval, and $\langle T \rangle$ is the observed stability duration.

25.5. Connection to CLIO

CLIO acts as a continuous limit of coherence updates. Increasing CLIO strength reduces effective ΔT, stabilizing projections. In the presence of irreducible obstruction $\mathcal{O}^* > 0$, the barrier $\Delta\mathcal{O}$ cannot be eliminated, so t_0 remains finite and a hard upper bound on stability persists.

> *Stability is not a primitive property of states, but the result of repeated suppression of admissible transitions.*

26. Gamma-Distributed Dwell Times from Obstruction Dynamics

Dwell times in bistable perception are gamma-distributed rather than exponential because escaping a perceptual state requires multiple partial advancements across the obstruction barrier, not a single event. Each step has an exponential waiting time; the sum follows a gamma. This section derives that distribution from first principles and reads its parameters back into the field geometry.

Empirical studies of bistable perception indicate that dwell times follow gamma or log-normal distributions rather than purely exponential ones. We show that gamma-distributed dwell times arise naturally from multi-step obstruction traversal.

26.1. Escape as a Multi-Step Process

A transition between perceptual states is not a single event but a sequence of k partial advancements across the obstruction barrier. Each advancement corresponds to a local alignment event:

$$\langle \mathbf{v}, \nabla\Phi \rangle > \theta.$$

Each step occurs stochastically with rate λ determined by local obstruction curvature $\nabla^2\mathcal{O}$, entropy S, and update interval ΔT.

26.2. Gamma Distribution

Each advancement has exponential waiting time $\lambda e^{-\lambda t}$. The total escape time $T = \sum_{i=1}^{k} t_i$ follows a gamma distribution:

$$P(T) = \frac{\lambda^k T^{k-1} e^{-\lambda T}}{(k-1)!}.$$

26.3. RSVP Interpretation of Parameters

- k: number of obstruction layers or required consistency alignments
- λ: effective alignment rate, modulated as $\lambda = \lambda_0 \cdot e^{-S} \cdot g(\Delta T)$

Higher entropy suppresses transitions (lower λ); more frequent updates reduce λ via Zeno stabilization; stronger obstruction increases k.

26.4. Special Cases

For $k = 1$ the distribution reduces to exponential, corresponding to shallow or weakly structured obstruction. For $k \gg 1$, by the central limit theorem, dwell times approach a Gaussian, corresponding to highly structured barriers requiring many incremental alignments.

26.5. Mean and Zeno Connection

The mean dwell time is $\langle T \rangle = k/\lambda$. Using $\lambda \sim \Delta T / t_0^2$:

$$\langle T \rangle \sim \frac{k\, t_0^2}{\Delta T},$$

recovering Zeno scaling as a special case of the gamma model with $k = 1$.

26.6. CLIO and Irreducible Complexity

CLIO modifies both parameters: it reduces k by removing unnecessary obstruction layers and increases λ by improving alignment efficiency. In the presence of irreducible obstruction, $k \geq k_{\min} > 0$, implying a lower bound on transition complexity.

26.7. Cognitive Interpretation

The parameters carry direct cognitive meaning:

- fast intuitive switches: low k, high λ
- deliberation: high k
- cognitive dissonance: high k, low λ
- expert intuition: low k, high λ

Observed switching statistics are a direct measurement of the geometry of obstruction in cognitive field space.

> *Perceptual switching is not a single escape event, but the accumulation of multiple constraint-satisfying transitions.*

27. Mixture Model: Random Obstruction Depth and Heavy Tails

Real dwell-time distributions are heavier-tailed than any single gamma can account for. This is because the depth of the obstruction barrier itself fluctuates—with

attention, fatigue, and context. Treating barrier depth as a random variable produces a gamma mixture, and when that depth is log-normally distributed (as multiplicative factors tend to be), heavy tails emerge naturally.

The gamma model of Section 26 treats k (the number of required obstruction crossings) as fixed. Empirically, dwell-time distributions often exhibit heavier tails than any single gamma can produce, and show log-normal-like behavior across subjects and conditions. This is naturally explained by treating k as itself a random variable, reflecting fluctuating obstruction geometry due to attention, fatigue, and context shifts.

27.1. Random Depth Model

Let k be drawn from a mixing distribution $p(k)$ over $\mathbb{Z}_{>0}$. The marginal dwell-time distribution is:

$$P(T) = \sum_{k=1}^{\infty} p(k) \cdot \frac{\lambda^k T^{k-1} e^{-\lambda T}}{(k-1)!}.$$

This is a gamma mixture, and its tail behavior is determined by $p(k)$.

27.2. Log-Normal Emergence

If k is log-normally distributed — which arises when obstruction depth is the product of many independent multiplicative factors (attention load, emotional modulation, contextual constraint) — then by the law of total variance, T inherits approximate log-normality for large mean k.

In RSVP terms, multiplicative factors on k correspond to:

- attention modulation of ΔT (Zeno stabilization rate),
- emotional weighting of Φ (salience potential depth),
- contextual shifts in S (entropy of accessible futures).

Each factor independently multiplies the effective obstruction depth seen by the system.

27.3. RSVP Interpretation

Fluctuating k corresponds to a time-varying obstruction landscape $\mathcal{O}(\pi, t)$. The curvature $\nabla^2_\pi \mathcal{O}$ at the current minimum determines the local effective k: deeper curvature requires more incremental advancements to escape. Attention and context shift the landscape, changing both k and λ simultaneously.

27.4. CLIO and Irreducible Floor

CLIO minimizes obstruction, reducing k toward $k_{\min}$. In the presence of irreducible obstruction, $k_{\min} > 0$ places a lower bound on transition complexity even after optimal repair. This implies:

- a minimum dwell time that cannot be shortened by coherence enforcement,
- a floor on switching variability that is structural rather than attentional,
- a diagnostic signature: systems near irreducible obstruction will show characteristic minimum dwell times resistant to Zeno stabilization.

Observed switching statistics are a direct measurement of the geometry of obstruction in cognitive field space — including its irreducible floor.

28. Katz-Admissibility and Spectral Obstruction in Receptor Operators

The receptor operator's activation condition—does a low-obstruction trajectory exist?—can be given a precise algebraic formulation using Katz-admissibility from algebraic geometry. A trajectory is admissible if and only if its cohomological obstruction class is trivial. This unifies the variational, topological, and cohomological accounts of perception into a single condition.

The receptor operator $\mathcal{R}_\Sigma$ defined in Section 22 activates when a trajectory γ has a low-obstruction representative in its homotopy class. The spectral interpretation of this condition — that the energy gap ΔE must lie in $\mathrm{Spec}(\Sigma)$ — has a precise algebraic formulation in terms of Katz-admissibility.

28.1. Katz-Admissibility

In the algebraic geometry setting, a local system or connection is Katz-admissible if its monodromy representation is compatible with the ramification structure of the base. Concretely, a trajectory's spectral data must be admissible with respect to the local ramification of the constraint surface Σ.

Translated into the receptor context: a trajectory γ is Katz-admissible with respect to Σ if and only if the energy transitions it requires lie within the ramification-allowed spectrum of Σ:

$$\Delta E(\gamma) \in \mathrm{Spec}_{\mathrm{Katz}}(\Sigma).$$

28.2. Connection to Obstruction Cohomology

The cohomological obstruction $[\omega] \in H^1(\mathcal{M}, \mathcal{F})$ introduced in the CLIO framework measures the failure of local sections to glue globally. At the receptor level, this obstruction is local: it measures whether the trajectory's spectral data can be consistently extended across the constraint surface.

A trajectory is Katz-admissible if and only if its associated cohomological obstruction class is trivial:

$$\gamma \text{ is Katz-admissible } \iff [\omega_\gamma] = 0 \in H^1(\Sigma, \mathcal{F}_\Sigma).$$

Receptor activation is therefore equivalent to triviality of the local obstruction class.

28.3. Interpretation

This unifies three levels of the theory:

1. *Variational*: activation $\iff \mathcal{E}_{\text{rec}}(\gamma) \approx 0$,
2. *Topological*: activation $\iff$ low-obstruction homotopy representative exists,
3. *Cohomological*: activation $\iff [\omega_\gamma] = 0$ in $H^1(\Sigma, \mathcal{F}_\Sigma)$.

All three are equivalent descriptions of the same event: a perturbation successfully integrating into the constraint structure.

28.4. Isotope Sensitivity as Ramification Change

Isotope substitution alters the vibrational spectrum without changing geometry. In Katz-admissibility terms, it changes the ramification type of the trajectory at the constraint surface: the substituted trajectory may fail to lie in $\text{Spec}_{\text{Katz}}(\Sigma)$ even though the unsubstituted one does. The cohomological obstruction class changes from trivial to nontrivial, and the receptor deactivates.

This provides the sharpest available formulation of isotope sensitivity: not a matter of frequency matching per se, but of ramification admissibility.

28.5. Generalization

The Katz-admissibility condition generalizes the receptor operator to any system where activation depends on spectral compatibility with a constraint structure. In cognitive systems, concept recognition is Katz-admissible when the incoming trajectory's spectral content is compatible with the constraint surface defined by the concept's internal structure. In symbolic systems, token acceptance is Katz-admissible when the token's embedding lies within the ramification-allowed spectrum of the grammar.

29. Operator Hierarchy: From Axiom to Realization

The paper introduces five distinct operator types at different levels of abstraction. This section makes their vertical connections explicit: a single axiom (consistency) propagates through geometry, dynamics, local testing, and observable statistics. Every result in the paper is a consequence of the axiom traversing these levels.

The preceding sections have introduced five distinct operator types. This section makes their hierarchical organization explicit, so the reader can see the full stack from foundational constraint to concrete implementation.

29.1. The Five Levels

Level 1: The Consistency Condition (Axiom).

$$\forall \gamma_i, \gamma_j \in \mathcal{M}, \quad \gamma_i \sim_{\mathcal{S}} \gamma_j \Rightarrow \gamma_i \sim \gamma_j.$$

This is the foundational constraint. Every subsequent level is either a consequence of this condition, a mechanism for enforcing it, or a diagnostic for measuring its violation.

Level 2: The Obstruction Functional (Geometry).

$$\mathcal{O}(\pi) = \sum_{i,j} \left| \mathbf{1}_{\gamma_i \sim \gamma_j} - \mathbf{1}_{\pi(\gamma_i)=\pi(\gamma_j)} \right|.$$

This converts the axiomatic condition into a measurable geometric quantity. Violation of Level 1 produces nonzero $\mathcal{O}$; the Obstruction–Distortion Proposition (Section 8) converts $\mathcal{O}$ into a lower bound on L_π.

Level 3: CLIO Dynamics (Repair).

$$\partial_t \pi = -\nabla_\pi \mathcal{O}(\pi).$$

This is the enforcement mechanism: a gradient flow in the space of projections that minimizes obstruction. Fixed points are locally consistent projections; irreducible fixed points (where $\mathcal{O}^* > 0$) are the formal signature of cognitive dissonance.

Level 4: The Local Variational Functional (Instance).

$$\mathcal{E}_{\text{rec}}(\gamma) = \mathcal{O}_{\text{local}}(\gamma) + \alpha\, \delta_{\text{local}}(\gamma)^2.$$

This instantiates the global framework at a local site (a receptor, a concept boundary, a token interface). Activation at a site is the existence of a γ achieving $\mathcal{E}_{\text{rec}} \approx 0$. Katz-admissibility (Section 28) provides the cohomological characterization of when such a γ exists.

Level 5: The Gamma Model (Statistics).

$$P(T) = \frac{\lambda^k T^{k-1} e^{-\lambda T}}{(k-1)!}, \qquad \lambda = \lambda_0 \cdot e^{-S} \cdot g(\Delta T).$$

This is the observable signature of the underlying field structure. The parameters (k, λ) encode the geometry of obstruction (depth and accessibility) in a form directly measurable from behavioral data. The mixture extension (Section 27) accounts for fluctuating geometry and produces heavy tails.

29.2. Vertical Connections

The levels are not independent modules but a connected derivation:

- Violation of Level 1 $\Rightarrow$ nonzero Level 2 $\Rightarrow$ increased L_π $\Rightarrow$ reduced stability bound $\lambda < (1 - L_F)/(L_B L_\pi)$.
- CLIO (Level 3) minimizes Level 2, driving the system toward Level 1 satisfaction.

- Irreducible obstruction sets a floor on Level 2, which propagates upward as a ceiling on λ and downward as a minimum k in Level 5.
- Katz-admissibility at Level 4 is equivalent to trivial obstruction cohomology, which is the local version of $\mathcal{O} = 0$ at Level 2.
- The Zeno relation $\langle T \rangle = t_0^2/\Delta T$ emerges from Level 3 dynamics applied to Level 2 geometry, and is a special case of Level 5 with $k = 1$.

29.3. Summary

The framework has one axiom (consistency), one geometry (distortion), one dynamics (CLIO), one local test (variational admissibility), and one observable signature (dwell-time distribution). Every result in this paper is a consequence of the axiom propagating through these levels.

30. Timeful Computation and the Rejection of Primitive Sequentiality

Fant's argument is that sequentiality is not primitive—dependency is. A clock imposes order from outside; a timeful system derives order from its own internal structure. This section reframes the entire architecture in those terms: Marine preserves dependency before windowing, RSVP evolves as a field rather than a state machine, TARTAN glues without serializing, and CLIO repairs as a wavefront rather than a procedure.

Fant's critique of clocked digital systems begins from the claim that sequentiality is not primitive. What is primitive is dependency: relations of enablement, inhibition, completion, and transition that possess innate concurrency. Sequential order is a derived imposition, typically enforced by a controller or clock that partitions concurrent behavior into externally managed intervals.

This matters for memory architecture because cognition is not naturally organized as a sequence of completed states. It is organized as overlapping wavefronts of partial activation, suppression, recurrence, and repair.

30.1. Marine Transduction as Preservation of Dependency

Standard windowed frequency analysis imposes a temporal interval before interpretation. In doing so, it converts concurrent micro-temporal structure into a stabilized spectral summary. The information discarded is not noise: it is the dependency structure of the signal itself.

The Marine Algorithm reverses this priority. Delta transduction treats local change as primary:

$$\Delta x_t = x_t - x_{t-1},$$

or more generally as an adaptive event condition:

$$|\Delta x_t| > \theta_t,$$

where θ_t is an adaptive threshold that prevents quantization artifacts from mimicking genuine micro-temporal fluctuation.

This preserves jitter not as nuisance but as the concurrent dependency relations of the incoming signal. Marine transduction is therefore not merely a preprocessing optimization. It is a structural refusal to collapse timeful dependency into clocked summary before those dependencies have been given the opportunity to compose.

30.2. RSVP as a Timeful Substrate

The RSVP state

$$X(x,t) = (\Phi, \mathbf{v}, S)$$

is intrinsically timeful. Here, Φ defines the field of potential differentiation, $\mathbf{v}$ defines propagating transition direction, and S defines configurational accessibility.

A cognitive process is therefore not a sequence of discrete states:

$$x_1 \to x_2 \to x_3,$$

but a propagation of constrained wavefronts through a field:

$$\partial_t X = \mathcal{D}(\Phi, \mathbf{v}, S).$$

Time is not imposed externally; it is expressed through field evolution. In Fant's terms, the vector field $\mathbf{v}$ corresponds precisely to the directed flow of wavefronts through a dependency network. RSVP provides the continuous substrate for what Fant's primitive behaviors do at the discrete level: negotiate concurrent dependency relations without requiring a clock to impose stability.

30.3. Rolling Coherence and Event Gating

A timeful architecture cannot rely on discrete batch boundaries. Burger's Recall Referee waits for a complete group of associative returns before enabling priority computation—a batch boundary that is clocked in spirit even if not in implementation.

A timeful replacement triggers arbitration by coherence accumulation. Define the normalized rolling coherence functional:

$$C(t) = \frac{1}{\tau} \int_{t-\tau}^{t} \langle \mathbf{v}(s), \nabla\Phi(s) \rangle \, e^{-S(s)} \, ds.$$

The normalization by τ renders $C(t)$ dimensionless and ensures that the coherence threshold η is a genuine intrinsic property of the system, independent of window length. As $\tau \to 0$, $C(t)$ approaches an instantaneous alignment measure; as $\tau \to \infty$, a time-averaged coherence. Both limits are interpretable without parameter retuning.

Arbitration becomes admissible when

$$C(t) > \eta.$$

This replaces clocked enable signals with coherence threshold crossings: the system commits when its own internal alignment justifies commitment, not when an external interval expires.

30.4. CLIO as a Consistency Wavefront

The procedural formulation of CLIO (Appendix A) operates by iteratively reducing obstruction one violating pair at a time. This description is useful for simulation and for establishing the existence of repair dynamics. It is not the deepest formulation.

A Fant-consistent CLIO is a field operator. Let $\mathcal{O}(\pi)$ denote the obstruction functional defined over the space of admissible projection maps $\pi : \mathcal{M} \to \mathcal{S}$. Then CLIO is the gradient flow:

$$\partial_t \pi = -\nabla_\pi \mathcal{O}(\pi),$$

where ∇_π denotes the functional gradient with respect to π, treating $\mathcal{O}$ as a functional on the space of quotient maps from $\mathcal{M}$ to $\mathcal{S}$. Repair is thus happening in the space of projections, not in the raw trace manifold $\mathcal{M}$.

Under this dynamics, local violations of the consistency condition generate correction gradients that propagate through the equivalence structure as repair wavefronts. Fixed points satisfy:

$$\nabla_\pi \mathcal{O}(\pi) = 0.$$

These are local minima, not guaranteed global minima. When $\mathcal{O}^* > 0$ at every fixed point, the irreducible obstruction of Section 33 is binding: no wavefront propagation can resolve the underlying structural inconsistency. Residual obstruction is structural information, not algorithmic failure.

The two formulations—procedural (Appendix A) and field-theoretic (this section)—are deliberately dual accounts of the same operator. The procedural version is falsifiable by simulation; the field-theoretic version is compatible with the continuous dynamics of the RSVP substrate.

30.5. TARTAN and Structural Concurrency

The TARTAN layer enforces sheaf-theoretic gluing conditions across overlapping memory tiles. In Fant's terms, this is the structural mechanism by which concurrent dependency relations in overlapping temporal regions are made to agree without imposing sequential order.

Where a clock enforces agreement by partitioning time into intervals, TARTAN enforces agreement by constraining overlapping regions to satisfy compatibility conditions on their intersections. Agreement is achieved through structure rather than temporal serialization. This is the continuous-domain analog of Fant's demand

for primitive behaviors that can negotiate concurrent dependency relations on their own expressional merits.

30.6. Conclusion

The architecture rejects primitive sequentiality at every major layer.

Marine preserves dependency before windowing. RSVP expresses cognition as field evolution rather than state transition. TARTAN enforces consistency across concurrent regions through structural gluing. CLIO propagates repair as a consistency wavefront in projection space. The interface layer discretizes only after coherence thresholds are crossed.

Sequence appears only at the boundary, as a projection for action or report. It is not the substrate. This is not an architectural preference but a consequence of what the consistency condition requires: any imposed sequential ordering that is not derivable from the dependency structure of the trace space is, in the formal sense of this paper, an obstruction.

31. The Projection Error and the Clip Economy

The formal machinery of the paper applies far beyond AI memory systems. This section traces the same obstruction structure through three scales: the projection error as an epistemological failure of state-based systems, the clip economy as its large-scale media instantiation, and biological life as the canonical example of a system that maintains its own constraint-closure loop against thermodynamic pressure.

The formal structure developed in this paper—obstruction, coarse-graining, consistency, repair—applies beyond AI memory architecture to the political economy of attention and media. This section connects those results to two phenomena: the *projection error* as a general epistemological failure of state-based systems, and the *clip economy* as its large-scale social instantiation.

31.1. The Projection Error

A system that records only local state transitions is blind to global topological structure. This is not a contingent limitation but a structural one: local linear transitions cannot distinguish processes whose difference is encoded in a loop rather than a path.

A biological example makes this precise. An exhaustive physical scan of a beating heart—chamber pressures, electrical conductivity, fluid dynamics, chemical concentrations—cannot distinguish between the heart's functional pumping of blood and the incidental sound it produces. Both are equal outputs of energy under the laws of physics. To distinguish them requires recognizing a *normative loop*: the blood feeds the muscle tissue that sustains the pump that circulates the blood. The heart maintains the organism and the organism maintains the heart.

Flattening this loop into a sequence of state transitions eliminates it. The property of being alive disappears from the data.

In RSVP terms, the projection error is the failure to preserve loop structure under coarse-graining. A projection $\pi : \mathcal{M} \to \mathcal{S}$ that collapses a closed trajectory γ (a loop) to a point or to an open path destroys the structural information that the loop encoded. This is a special case of false coherence: distinct traces—the closed loop and its collapsed image—are identified in $\mathcal{S}$ despite being inequivalent in $\mathcal{M}$.

The normative loop is precisely the structure that the Consistency Condition (Theorem 6.1) is designed to protect. A valid quotient map must preserve recurrence structure, which includes the recurrence of a trajectory with itself—the defining property of a loop. Invalid compression does not merely forget; it actively destroys the topological invariants that distinguish living from dead, functional from incidental, signal from noise.

31.2. Optimization for Recoverability

The projection error generates a design imperative: systems should optimize not for compression but for *structural recoverability*—the preservation of enough information to reconstruct the global context from local fragments.

This is the engineering translation of the Consistency Condition. A compression that satisfies the condition is recoverable in the relevant sense: the equivalence classes of the quotient are unions of equivalence classes of the original, so the original structure can be recovered up to the granularity of the quotient. A compression that violates the condition is not recoverable: the classes have been defined by external fiat, and no inverse map can reconstruct what was never derived from the trace space.

Recoverability is therefore not a post-hoc desideratum but a constraint that falls directly out of the consistency requirement. Systems that fail to satisfy it are not merely lossy; they are structurally hallucinating.

31.3. The Clip Economy as Distributed Projection Error

At scale, the projection error manifests as what may be called the *clip economy*: an information environment optimized for the compression and circulation of atomic, decontextualized fragments—clips, excerpts, highlights, thumbnails—stripped of the trajectory context that gave them meaning.

The formal structure is the same. Each clip is a local state transition: a fragment of a longer process, projected out of the global loop it belonged to. The recommendation and feed systems that distribute these fragments are themselves coarse-graining operators—they sort by engagement signal (a local metric) rather than by structural coherence (a global property). The result is a system with high obstruction: equivalence classes in the distributed representation space that have no valid preimage in the trajectory space of the original content.

In RSVP terms, the feed is a projection $\pi_{\text{feed}} : \mathcal{M}_{\text{content}} \to \mathcal{S}_{\text{feed}}$ where the equivalence classes of $\mathcal{S}_{\text{feed}}$ are defined by engagement metrics rather than recurrence structure. This generically violates the Consistency Condition. The cognitive experience of consuming such a feed is precisely the phenomenology of high obstruction: fragments that cannot be integrated, patterns that cannot cohere, structure that cannot be recovered.

The stability theorem (Theorem 7.1) applies directly. A cognitive system that relies heavily on a high-obstruction external representation (π_{feed}) inherits the distortion of that projection. The effective L_π of the feed enters the stability bound on the system's self-referential dynamics. High feed obstruction reduces the safe gain of the cognitive subconscious layer—the system must rely less on its own compressed history precisely when the external information environment is most incoherent.

31.4. Admissible Fragments and Reconstruction Pressure

The corrective requirement is that fragments carry sufficient structural information to license their reintegration. An *admissible fragment* in the sense of this paper is one for which a valid lift exists: a distribution μ over the trace space $\mathcal{M}$ such that $\pi_*(\mu) = \delta_s$ and μ is supported on recurrence-consistent traces (cf. Section 18, Interface Admissibility Condition).

An admissible fragment carries what might be called *reconstruction pressure*: it not only conveys local content but also constrains the global context into which it must be placed. It is the difference between a puzzle piece that has information about the image it belongs to and one that does not.

The engineering implication is that media and information systems should be evaluated not only by compression efficiency or engagement signal, but by the obstruction score of the coarse-graining they implement. Systems with low obstruction preserve context; systems with high obstruction destroy it, and the destruction is not recoverable downstream.

31.5. Life as a Constraint-Closure Loop

The projection error discussion above identifies a minimal formal definition of biological life that emerges from the RSVP framework: a living system is one that maintains a stable, self-referential constraint-closure loop against the pressure of thermodynamic dissipation.

More precisely, a living region $R \subset X$ is characterized by the existence of a stable endomorphism $\phi : R \to R$ such that:

1. ϕ is not the identity (the system actively maintains itself, not just persists),
2. ϕ is constraint-preserving (the conditions of the system's operation are reproduced, not merely inherited),
3. ϕ is stable under small perturbations (the loop is robust, not fragile).

A machine has its boundary conditions imposed externally by an engineer. A living

cell continuously rebuilds its own membrane—the wall that protects the chemical reactions that build the wall. It manufactures its own constraints rather than receiving them.

In sheaf-theoretic terms, the health of a nested biological system (cell → organ → organism → ecosystem) depends on the gluing conditions between levels. Full admissibility corresponds to health: local constraint loops preserve internal autonomy while contributing coherently to higher-level loops. Partial admissibility corresponds to compensated pathology: some loops degrade and gluing conditions weaken, but the global section persists. Zero admissibility corresponds to collapse: the restriction map between levels fails entirely.

Cancer is zero admissibility at the cellular level. The tumor cell's internal loop is highly optimized—it replicates effectively—but it decouples from the organism's global section. It treats the host as an external resource rather than a shared constraint structure. The gluing condition fails; the cell forgets the whole. This is formally identical to the fragment in the clip economy that maximizes local engagement by destroying global context: high local Φ, zero contribution to the global section.

Medicine, in this framework, is re-gluing: restoring the topological conditions under which a global section—a unified, coherent organism—can exist. CLIO at the biological scale.

Summary. The projection error, the clip economy, and biological pathology are three scales of the same formal phenomenon: the destruction of constraint-closure structure by coarse-graining operations that violate the Consistency Condition. The corrective in all three cases is the same: preserve recurrence structure, enforce gluing conditions, and interpret residual obstruction as structural information rather than noise.

32. Natural Implementation Patterns and Their Structural Consequences

Theory is most useful when it can name the specific ways that reasonable engineering decisions go wrong. This section examines three patterns that arise naturally when building wave-based memory systems—not from carelessness, but from sound intuitions applied to the wrong invariants. Each pattern is described, its structural consequence identified, and the correct replacement given.

The preceding formal apparatus identifies structural requirements that are easy to violate in practice—not through carelessness, but because the most natural engineering decisions at each stage tend to produce exactly the failure modes the theory predicts. This section examines three such patterns. Each is a reasonable first instinct; each introduces a specific, identifiable obstruction. Understanding why these patterns arise naturally is as important as understanding why they fail structurally.

32.1. Operationalizing Holonomy as a Frequency Threshold

When building a monitoring layer intended to detect path-dependent inconsistency, a natural implementation strategy is to check whether a memory's pulse history falls within an expected oscillatory band—for instance, verifying that long-duration traces cycle within a 0.73 Hz resonance window. This is measurable, efficiently computable, and intuitively connected to the temporal structure of the system.

The difficulty is structural. A frequency threshold is a local spectral condition: it checks whether a given trace occupies an expected region of power spectrum. Holonomy, as defined in Section 11, is a path-dependent quantity: it measures whether transporting an equivalence class around a closed loop in trace space returns it to its starting point. These are not the same condition, and they fail in orthogonal ways.

A trace can satisfy the frequency threshold while accumulating nonzero holonomy: the oscillation is present, but the equivalence structure drifts across the loop. Conversely, a trace may exhibit holonomy defect $H(\gamma) = 0$ while temporarily falling outside the expected spectral band—a transient fluctuation that the threshold misidentifies as structural failure.

More precisely: the holonomy defect $H(\gamma) = d(\pi_1(\gamma(0)), \pi_0(\gamma(0)))$ is a property of the projection map π along a path, not a property of the signal's power at a given frequency. Replacing the former with the latter conflates the geometry of the projection space with the spectral content of the input. The monitoring layer then becomes sensitive to the wrong invariant, raising false positives on benign fluctuations and missing genuine inconsistency accumulation.

The structurally correct implementation detects holonomy by tracking how the induced equivalence relation $\sim_{\mathcal{S}}$ evolves as the system traverses loops in trace space—specifically, whether a class that was identified at the start of a loop is re-identified at its end. This requires maintaining projection history across trajectories, not just a running spectral summary. The additional bookkeeping is not optional; it is the measurement.

32.2. Mapping Valence onto a Fixed Emotional Submanifold

A second natural pattern arises when implementing emotional modulation. The intuition is sound: emotional states should influence the geometry of retrieval, and it is natural to represent this by projecting the system's wave amplitude onto a submanifold whose coordinates are organized by valence—positive affect toward one region, negative affect toward another.

This is the frequency-mapping failure (Section 2) expressed in geometric language. The valence axis is an externally specified coordinate: it does not derive from the recurrence structure of the raw trace space $\mathcal{M}$, but from a prior categorization of emotional states imposed at design time. The submanifold M_e is defined by fiat rather than discovered from trace dynamics.

The Consistency Condition (Theorem 6.1) requires that the fibers of any projection $\pi : \mathcal{M} \to \mathcal{S}$ be unions of equivalence classes already present in $\mathcal{M}$. A valence-based submanifold projection fails this condition whenever the valence partition does not align with the recurrence structure of the traces. Two traces that are recurrently similar but emotionally labeled differently will be separated into distinct submanifold regions; two traces that are recurrently distinct but emotionally similar will be merged. Both constitute obstruction.

The instability consequence follows from Section 8: each such misalignment increases the bi-Lipschitz defect δ_π, tightening the stability bound on admissible subconscious gain. A system with many valence-misaligned traces is not merely imprecise—it has structurally reduced the safe operating range of its own emotional modulation layer.

The corrective is not to abandon emotional geometry but to derive it. The submanifold structure should emerge from the clustering of recurrent traces under the induced equivalence relation $\sim$, with valence appearing as a secondary labeling of the resulting clusters rather than as their defining axis. The geometry comes first; the interpretation follows.

32.3. Batch-Gating Recall Arbitration

A third pattern concerns the arbitration of competing memory retrievals. When multiple associative returns arrive in parallel, a natural implementation waits for a complete group of returns before computing priority—collecting all candidates, then resolving conflicts by comparing their coherence scores. This is clean, deterministic, and easy to test.

The difficulty is that this strategy imposes a batch boundary: the system holds execution until an externally defined collection is complete, then commits. This is a clock in the sense of Section 30, even when it is not implemented as a literal timer. The commitment point is determined by the size of the candidate set, not by the internal coherence dynamics of the system.

The consequences are two. First, the batch boundary introduces a latency that is independent of the system's actual state of coherence: a highly consistent set of returns is held as long as a highly inconsistent one. The system cannot distinguish between "all candidates have arrived and agree" and "all candidates have arrived and conflict"; it only detects that the group is complete.

Second, and more fundamentally, the batch model requires that all relevant returns arrive before arbitration begins. In a timeful architecture (Section 30), retrieval is not a discrete event but a wavefront: returns arrive with varying latency as the projection dynamics propagate through the trace space. A batch boundary that fires when the first n returns arrive will systematically exclude late-arriving but structurally significant traces—precisely those that require the most traversal to surface, and therefore carry the most distinctive recurrence information.

The structurally correct replacement is coherence-threshold gating. The normalized

rolling coherence functional

$$C(t) = \frac{1}{\tau} \int_{t-\tau}^{t} \langle \mathbf{v}(s), \nabla\Phi(s) \rangle \, e^{-S(s)} \, ds$$

rises continuously as compatible returns accumulate. Arbitration becomes admissible when $C(t) > \eta$: the system commits when its own internal alignment justifies commitment, not when a predefined collection size is reached. This preserves the timeful structure of retrieval while providing a well-defined trigger for resolution.

32.4. The Common Structure

All three patterns share a single underlying form: a global, path-dependent, or dynamically emergent property is replaced by a local, static, or externally specified proxy. Holonomy becomes a frequency band; emergent emotional geometry becomes a fixed valence axis; coherence-driven commitment becomes a batch count.

In each case the proxy is cheaper to compute and easier to specify. In each case it measures a different invariant than the one required—one that does not survive the structural conditions the theory imposes.

> *The natural implementation is not wrong because it is naive. It is wrong because it answers a different question than the one the system is actually asking.*

Identifying these patterns early—before they are load-bearing in a larger architecture—is the practical value of maintaining a formal account of what valid compression requires. The three correctives described above are not exotic: they require tracking projection history rather than spectral summary, deriving geometric structure rather than imposing it, and gating on coherence rather than collection size. The additional complexity in each case is exactly the complexity the formal constraint demands.

33. Conclusion

Mem|8, when stripped of its representational artifacts, reveals a fundamentally sound direction: recursive memory as self-coarse-graining. The wave-based architecture, emotional modulation, and pattern functional $P(t)$ are all structurally motivated. What the system lacks is a formal account of when its own compression is valid.

The framework developed here provides that account. Valid subconscious structure must be a coarsening of recurrence already present in the raw trace space. Violations produce obstruction. Obstruction induces metric distortion. Distortion destabilizes self-referential dynamics. Repair reduces obstruction—but only to a limit.

Irreducible Obstruction and Gain Limitation

The preceding analysis implicitly assumes that $\mathcal{O}(\pi)$ can be reduced arbitrarily by repair. However, this need not be the case.

Let $\mathcal{O}^* > 0$ denote an irreducible obstruction: a lower bound such that no admissible repair operator, including CLIO, can produce a projection π with $\mathcal{O}(\pi) < \mathcal{O}^*$.

By Proposition 8.1, $\mathcal{O}^*$ induces a corresponding lower bound on geometric distortion:

$$L_\pi \geq f\left(\frac{\mathcal{O}^*}{|\mathcal{M}|^2}\right) =: L_\pi^{\min}.$$

Substituting into Theorem 7.1:

$$L_F + \lambda L_B L_\pi^{\min} < 1,$$

we obtain a hard constraint on admissible subconscious gain:

$$\lambda < \frac{1 - L_F}{L_B L_\pi^{\min}}.$$

This ceiling is not algorithmic but structural. No improvement in the repair operator can raise it. It reflects incompatibility in the underlying trace space itself.

A system with irreducible obstruction must therefore choose between reducing λ—diminishing reliance on subconscious compression—or accepting violation of the stability condition. There is no third option.

Final Remark. Cognitive dissonance is not merely a representational inconsistency. It is a dynamical constraint: it reduces the safe operating range of self-referential influence. The subconscious must whisper more quietly precisely when its internal coherence is least assured.

This is not a limitation that better architecture can engineer away. It is the shape of the problem.

Projection error, injection error, and holonomy are not independent. They form a hierarchy: pairwise inconsistency (obstruction) produces geometric distortion, which accumulates along trajectories as holonomy. Stability constraints propagate this structure into dynamics, limiting admissible self-reference. A system that is locally consistent but globally inconsistent will not fail immediately. It will fail historically.

The subconscious does not compute, select, or interpret. It minimizes. When the minimization is well-posed, coherence emerges as a byproduct. When it is not, inconsistency propagates geometrically and temporally, and no amount of local correction can restore global stability.

The subconscious must whisper, because it is solving a constrained problem. When the constraints tighten, the whisper is not a choice—it is a necessity.

References

[1] L. Turin, A spectroscopic mechanism for primary olfactory reception, *Chemical Senses*, 21(6):773–791, 1996.

[2] L. Turin, *The Secret of Scent: Adventures in Perfume and the Science of Smell*, Harper Perennial, 2014.

[3] N. Turok, *The Universe Within: From Quantum to Cosmos*, House of Anansi Press, 2018.

[4] L. Boyle, K. Finn, N. Turok, CPT-Symmetric Universe, *Physical Review Letters*, 121(25):251301, 2018.

[5] N. Turok, A New View of the Universe, *Nature Astronomy*, 3:875–876, 2019.

[6] N. Turok, Beyond Inflation: A New Paradigm for the Early Universe, *Proceedings of the Royal Society A*, 478(2265), 2022.

[7] J. C. Brookes, F. Hartoutsiou, A. P. Horsfield, A. M. Stoneham, Could humans recognize odor by phonon assisted tunneling? *Physical Review Letters*, 98(3):038101, 2007.

[8] L. B. Vosshall, Olfaction in Drosophila, *Current Opinion in Neurobiology*, 10(4):498–503, 2000.

[9] R. Moreno-Bote, J. Rinzel, N. Rubin, Noise-induced alternations in an attractor network model of perceptual bistability, *Journal of Neurophysiology*, 98(3):1125–1139, 2007.

[10] J. Brascamp, R. Blake, Inattention abolishes binocular rivalry: perceptual evidence, *Journal of Vision*, 8(5):1–11, 2008.

[11] W. J. M. Levelt, *On Binocular Rivalry*, Van Gorcum, 1965.

[12] H. Haken, *Synergetics: An Introduction*, Springer, 1983.

[13] P. Ashwin, S. Coombes, R. Nicks, Mathematical frameworks for oscillatory network dynamics in neuroscience, *Journal of Mathematical Neuroscience*, 6(2), 2016.

[14] C. Gardiner, *Stochastic Methods: A Handbook for the Natural and Social Sciences*, Springer, 2009.

[15] H. A. Kramers, Brownian motion in a field of force and the diffusion model of chemical reactions, *Physica*, 7(4):284–304, 1940.

[16] N. G. van Kampen, *Stochastic Processes in Physics and Chemistry*, North-Holland, 2007.

[17] K. M. Fant, *Logics of Time and Computation*, PhD Thesis, Chalmers University of Technology, 1994.

[18] R. Milner, *Communication and Concurrency*, Prentice Hall, 1989.

[19] S. Abramsky, A. Jung, Domain theory, in *Handbook of Logic in Computer Science*, Vol. 3, 1994.

[20] S. Mac Lane, *Categories for the Working Mathematician*, Springer, 1998.

[21] N. M. Katz, *Rigid Local Systems*, Princeton University Press, 1996.

[22] A. Grothendieck, *Revêtements Étales et Groupe Fondamental (SGA 1)*, Springer Lecture Notes in Mathematics, 1971.

[23] G. E. Bredon, *Sheaf Theory*, Springer, 1997.

[24] M. Gromov, Structures métriques pour les variétés riemanniennes, CEDIC, 1981.

[25] C. Villani, *Optimal Transport: Old and New*, Springer, 2009.

[26] T. Cover, J. Thomas, *Elements of Information Theory*, Wiley, 2006.

[27] C. E. Shannon, A mathematical theory of communication, *Bell System Technical Journal*, 27(3):379–423, 1948.

[28] K. Friston, The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138, 2010.

[29] P. C. Bressloff, *Stochastic Processes in Cell Biology*, Springer, 2014.

[30] M. Wilson, *Wandering Significance*, Oxford University Press, 2006.

[31] M. Wilson, *Physics Avoidance*, Oxford University Press, 2017.

A. Numerical Simulation of Obstruction and CLIO Repair

This appendix describes a minimal simulation demonstrating how subconscious hallucination arises from invalid coarse-graining and how CLIO repair reduces obstruction.

A.1. Trace Space

Let the raw trace space consist of points

$$\mathcal{M} = \{\gamma_1, \gamma_2, \ldots, \gamma_n\}$$

embedded in $\mathbb{R}^2$. A recurrence relation is defined by distance:

$$\gamma_i \sim \gamma_j \iff \|\gamma_i - \gamma_j\| < \varepsilon.$$

A.2. Obstruction Score

Define:

$$\mathcal{O}(\pi) = \sum_{i,j} \left| \mathbf{1}_{\gamma_i \sim \gamma_j} - \mathbf{1}_{\pi(\gamma_i)=\pi(\gamma_j)} \right|.$$

A consistent projection satisfies $\mathcal{O}(\pi) = 0$.

A.3. CLIO Repair Algorithm

CLIO proceeds via local repair moves:

1. Identify a violating pair (γ_i, γ_j).
2. If $\gamma_i \sim \gamma_j$ but $\pi(\gamma_i) \neq \pi(\gamma_j)$: merge their subconscious classes (resolve false conflict).
3. If $\gamma_i \not\sim \gamma_j$ but $\pi(\gamma_i) = \pi(\gamma_j)$: split the class if a valid partition exists (resolve false coherence).
4. Accept the move if it reduces $\mathcal{O}$.
5. Repeat until $\mathcal{O}$ does not decrease.

CLIO is greedy descent on $\mathcal{O}(\pi)$ and is not guaranteed to reach a global minimum. Residual obstruction at termination is interpreted as irreducible—the formal signature of cognitive dissonance (see Section 17).

A.4. Pseudocode

```
Input:
  traces gamma[1..n], epsilon, initial_projection pi

Compute raw equivalence:
  raw_same[i,j] = (distance(gamma[i], gamma[j]) < epsilon)

Compute subconscious equivalence:
  sub_same[i,j] = (pi[i] == pi[j])

Obstruction:
  O = sum over i,j of |raw_same[i,j] - sub_same[i,j]|

CLIO repair loop:
  while O > 0:
    find violating pair (i,j)

    if raw_same[i,j] and not sub_same[i,j]:
      merge pi[i] and pi[j]          # resolve false conflict

    if not raw_same[i,j] and sub_same[i,j]:
      split one unsupported assignment # resolve false coherence

    recompute O

    if O does not decrease:
      flag irreducible obstruction
      stop
```

```
Output:
  repaired projection pi, residual 0, obstruction flag
```

A.5. Interpretation

Before CLIO repair, the subconscious layer may contain equivalence classes with no valid preimage in the raw trace space. These appear as hallucinated attractors. After CLIO repair, subconscious classes become valid quotients of recurrence-supported trace classes.

Hallucination is not merely the production of false content. It is the formation of invalid equivalence structure. CLIO repairs hallucination by restoring compatibility between the raw recurrence relation and the subconscious quotient.

If the exit condition is reached with $\mathcal{O} > 0$, this does not indicate failure of the repair process. It indicates that the raw traces themselves contain incompatible structure—and that the hard ceiling on λ derived in Section 33 is binding.