MEM|8

Memory Before Representation

Recoverability, Continuation, and the Geometry of Persistent Structure

Flyxion

Independent Research

Canada

Draft — June 27, 2026

A memory is not a preserved past.

A memory is the continued ability
to regenerate distinctions from transformation.

Contents

Preface

This monograph develops the MEM|8 framework from first principles. The central claim is simple but has far-reaching consequences: *memory is not storage; memory is recoverable continuation.*

No physical system preserves a past state unchanged. Every substrate undergoes continual transformation at every timescale. The question is therefore never whether a state remains, but whether the distinctions that once were present can be reconstructed from the state that now exists.

That shift from storage to recoverability is not a minor revision of memory theory. It is a change of ontological category. Traditional accounts place representation first and derive memory as its persistence. MEM|8 reverses the order: recoverable distinctions are primary, and representations emerge from stable patterns of recoverability.

The formal architecture rests on three empirical anchors drawn from contemporary physics, neuroscience, and cell biology. The topological acoustic synapse (TAS) provides a physical demonstration that a small substrate can generate a large high-dimensional space of reconstructive witnesses. Intrinsic macroscale oscillatory modes in fMRI data provide the dynamical picture: memory states are resonances of a latent modal basis, and recall is thresholded excitation rather than lookup. Synaptic vesicle trafficking provides the maintenance picture: persistence is not component preservation but flux balance, a continuous renewal that conserves reconstructive organization while replacing every material carrier.

Together these three anchors yield the MEM|8 synthesis:

$$\boxed{\text{Memory} = \text{recoverable resonant structure maintained by flux.}}$$

The monograph is arranged in eight parts. Parts I–III build the mathematical foundations. Parts IV–V treat ecological and social memory. Part VI connects MEM|8 to computation. Part VII extends the theory to civilization scale. Part VIII addresses the foundational inversions: representation, identity, and time all emerge from memory rather than preceding it.

Mathematical prerequisites are a first course in real analysis, linear algebra, and ordinary differential equations. Familiarity with Hilbert spaces is helpful but not assumed; the necessary functional analysis is developed inline.

Flyxion
Canada, June 27, 2026

Part I

The Failure of Storage Metaphors

Chapter 1

The Storage Illusion

1.1 Why Storage Cannot Be Primitive

The dominant picture of memory inherited from computer science treats memory as a container: states are written in, preserved, and read out. This picture is useful as an engineering abstraction, but it misrepresents the physical situation.

Proposition 1.1 (No Physical Substrate Is Stably Inert)**.** *No physical substrate under realistic conditions preserves an exact microstate indefinitely. More precisely, for any substrate S at finite temperature $T > 0$, the probability that $S(t) = S(t_0)$ exactly decays to zero as $t \to \infty$.*

Proof. Under Hamiltonian dynamics a substrate may return arbitrarily close to an earlier state (Poincaré recurrence), but exact microstate equality $S(t) = S(t_0)$ for $t > t_0$ is measure-zero in phase space. At finite temperature, coupling to a thermal bath introduces stochastic forcing $\eta(t)$; the Fokker–Planck equation for the state distribution has no point-mass stationary solution at $S(t_0)$, so the probability of exact preservation decays strictly with time. Under dissipative dynamics without thermal noise a fixed point may exist, but biological and cognitive substrates operate far from equilibrium and are therefore excluded from this case. In all realistic conditions, exact microstate preservation is either measure-zero or dynamically unstable under perturbation. □

Remark 1.2. The proposition is not that states never recur approximately — they may, as harmonic oscillators demonstrate. The point is that exact preservation is not a stable property: any perturbation, however small, eventually destroys it. Memory therefore cannot be grounded in exact preservation; it must be grounded in something more robust.

The consequence is immediate: if memory required exact preservation of a past state, no physical system could be said to remember anything. Since manifestly some systems do remember, the storage account must be wrong.

Definition 1.3 (Storage Account)**.** A *storage account* of memory holds that a system S remembers x at time t if and only if there exists a subsystem $\sigma \subseteq S$ such that

$$\sigma(t) = \sigma(t_0) \quad \text{and} \quad \sigma(t_0) \text{ encodes } x.$$

Proposition 1.1 shows that the antecedent of the storage account is never physically satisfied. The account must therefore be abandoned or radically revised.

1.2 What Survives Change

Something clearly does survive. A person remembers their childhood; a city retains the imprint of its construction; a document preserves the intent of its author. The question is what kind of thing survives when no state literally persists.

The answer proposed here is a *capacity*: the capacity to regenerate, from the current state, the distinctions that were once present.

Definition 1.4 (Reconstructive Capacity)**.** A system S has *reconstructive capacity* for a distinction δ at time t if there exists a physically realizable process $\mathcal{R}$ such that

$$\mathcal{R}(S(t)) \;\approx_\varepsilon\; \delta,$$

where $\approx_\varepsilon$ denotes approximate recovery within tolerance $\varepsilon > 0$.

This shifts the question from "what is stored?" to "what can be reconstructed?" The shift is more than terminological. It opens the way to a theory in which memory is graded, ecological, and dynamical rather than binary, internal, and static.

Chapter 2

Memory as Reconstruction

2.1 Reconstruction Operators

Definition 2.1 (Reconstruction Operator)**.** Let $\mathcal{H}$ be the space of admissible states of a system and let $\mathcal{Q}$ be a space of distinctions (queries, labels, categories). A *reconstruction operator* is a map

$$R_t : \mathcal{H} \longrightarrow \mathcal{Q},$$

indexed by current time t, whose intended action is to recover, from the present state $S(t) \in \mathcal{H}$, the distinction that was present at the target time t_0.

Definition 2.2 (Memory Strength)**.** Given a metric $d_{\mathcal{Q}}$ on the distinction space, the *memory strength* of a system for distinction $q_0 = R_{t_0}(S(t_0))$ at current time t is

$$M(t_0, t) \;=\; 1 - d_{\mathcal{Q}}\big(R_t(S(t)),\; q_0\big).$$

Perfect memory corresponds to $M = 1$; complete forgetting to $M = 0$ (or below some operational threshold).

Remark 2.3. Memory strength is not a property of the substrate alone. It depends jointly on the current state $S(t)$, the reconstruction operator R_t, and the target distinction q_0. This three-way dependence is absent from storage accounts, which locate memory entirely in the substrate.

2.2 Admissible Recovery

Not every reconstruction is admissible. A system that produces an arbitrary output trivially achieves "reconstruction" in a vacuous sense. Admissibility constrains which outputs count.

Definition 2.4 (Admissibility Metric)**.** An *admissibility metric* on $\mathcal{Q}$ is a metric $d_{\mathcal{Q}}$ such that two distinctions q, q' are considered the same recovered memory if and only if $d_{\mathcal{Q}}(q, q') \leq \varepsilon$ for a contextually fixed tolerance ε.

Proposition 2.5 (Memory Is a Pseudometric)**.** *Define*

$$\rho(t_0, t) = d_{\mathcal{Q}}\big(R_t(S(t)),\ R_{t_0}(S(t_0))\big).$$

Then ρ is a pseudometric on the space of (t_0, t) pairs.

Proof.

1. *Non-negativity.* $\rho \geq 0$ because $d_{\mathcal{Q}}$ is a metric.
2. *Identity.* $\rho(t_0, t_0) = 0$ because $R_{t_0}(S(t_0)) = R_{t_0}(S(t_0))$.
3. *Symmetry.* $\rho(t_0, t) = \rho(t, t_0)$ by symmetry of $d_{\mathcal{Q}}$.
4. *Triangle inequality.*

$$\begin{aligned}\rho(t_0, t_2) &= d_{\mathcal{Q}}(R_{t_2}(S(t_2)),\ R_{t_0}(S(t_0)))\\ &\leq d_{\mathcal{Q}}(R_{t_2}(S(t_2)),\ R_{t_1}(S(t_1))) + d_{\mathcal{Q}}(R_{t_1}(S(t_1)),\ R_{t_0}(S(t_0)))\\ &= \rho(t_1, t_2) + \rho(t_0, t_1).\end{aligned}$$

Therefore ρ satisfies all pseudometric axioms. (It fails to be a metric because two states with identical reconstructive outputs may differ as physical states.) $\square$

Chapter 3

The Epistemic Necessity of Memory

Definition 3.1 (Distinction Capacity)**.** Write $\delta_R(t)$ for the number of distinct outputs a reconstruction operator R_t can discriminate among at time t.

Theorem 3.2 (Collapse of Temporal Comparison)**.** *If $\delta_R(t) = 0$ for all t, then no temporal comparison, learning, or reference is possible.*

Proof. Temporal comparison requires distinguishing the state at t_0 from the state at t_1. If $\delta_R = 0$, every reconstruction is identical, so no pair (t_0, t_1) is distinguishable by any query. Learning requires detecting a change between an initial state and a later state; if no states are distinguishable, no change is detectable. Reference requires identifying a past target; with $\delta_R = 0$ every past target maps to the same output. All three capacities therefore collapse to zero. □

Corollary 3.3. *Cognition, science, and language are only possible in systems that maintain nonzero reconstructive capacity. Memory, in the sense of Definition 1.4, is a necessary condition for any epistemic activity.*

Part II

Event Memory

Chapter 4

Events Before Objects

4.1 The Primacy of Events

The storage account places objects first: there are things that exist, and memory stores facts about them. MEM|8 inverts this order. Objects are not given; they are reconstructed. What is given are events, and objects are the stable patterns that emerge from reconstruction across event sequences.

Definition 4.1 (Event Sequence)**.** An *event sequence* is a finite or countably infinite ordered collection

$$E = (e_1, e_2, \ldots, e_n, \ldots)$$

where each e_i belongs to an event space $\mathcal{E}$.

Definition 4.2 (Object as Compressed History)**.** An *object identity* $\mathcal{O}$ relative to an event sequence E and reconstruction operator R is an equivalence class

$$\mathcal{O} = \{e_i \in E : R(H_i) \approx_\varepsilon R(H_j) \text{ for all } i, j\}$$

where $H_i = \{e_k : k \leq i\}$ is the history up to event i.

In words: an object is whatever a reconstruction operator assigns consistently the same answer across a stretch of the event sequence. Object identity is derived, not primitive.

Chapter 5

Histories as Computational Structures

Definition 5.1 (Historical State)**.** The *historical state* at time t is

$$H_t = \{e_i \mid i \leq t\},$$

the complete record of events up to and including t.

Definition 5.2 (Compression Ratio)**.** Let $K(H_t)$ denote the Kolmogorov complexity of H_t. The *compression ratio* of the present state $S(t)$ as a witness to H_t is

$$\chi(t) = \frac{K(H_t)}{|S(t)|},$$

where $|S(t)|$ is the description length of the current physical state. A system is an *efficient witness* of its history when $\chi(t) \gg 1$.

Proposition 5.3 (Present State as Compressed Witness)**.** *The present physical state of a system is a lossy compression of its historical trajectory. Memory is the capacity to partially decompress this compression on demand.*

Sketch. The present state $S(t)$ is a function of all past states via the dynamics: $S(t) = F(S(0), e_1, \ldots, e_t)$. The map F is generally not injective; many distinct histories H_t can produce the same $S(t)$. Therefore $S(t)$ is a compression of H_t, and R_t is the corresponding decompression applied selectively to the query at hand. $\square$

Chapter 6

The Event-Log Ontology

Definition 6.1 (Event-Log System)**.** A system Σ is an *event-log system* if its current state can be interpreted as a compressed summary of an event log and if that summary supports at least partial reconstruction of historical distinctions.

Example 6.2. The following are event-log systems under this definition:

1. *Biological cells.* The epigenetic state summarizes developmental and environmental event history.

2. *Brains.* Synaptic weights compress a lifetime of sensorimotor events.

3. *Documents.* Text preserves the communicative acts that produced it.

4. *Institutions.* Policies, procedures, and personnel encode accumulated decisions.

5. *Languages.* Phonological and grammatical structures reflect historical contact events across centuries.

6. *Ecosystems.* Species composition records evolutionary and climatic history.

The event-log ontology unifies these diverse systems under a single memory-theoretic description. Memory is not a special faculty of minds; it is a property of any system whose present state is a compressed witness of past events.

Part III

The Mathematical Architecture of Recall

Chapter 7

Memory as Recoverable Modal Structure

7.1 The Memory Field

Definition 7.1 (Memory Field). Let Ω be a cognitive or physical domain and let

$$\mathcal{H} = L^2(\Omega; \mathbb{C})$$

be the Hilbert space of admissible memory states. A *memory field* is a time-indexed section

$$M : \mathbb{R}_{\geq 0} \longrightarrow \mathcal{H}, \qquad t \mapsto M(t),$$

whose instantaneous state admits a modal expansion

$$M(t) = \sum_{i \in I} a_i(t)\, \psi_i,$$

where $\{\psi_i\}_{i \in I}$ is an orthonormal family of *reconstructive modes* and $a_i(t) \in \mathbb{C}$ are time-dependent excitation coefficients.

The modes $\{\psi_i\}$ are latent structures of the memory system; they are determined by the physical or cognitive substrate, not by the content of any individual memory. Individual memories are patterns of excitation over this fixed modal basis.

7.2 Memory Is Not Substrate Identity

Definition 7.2 (Memory Equivalence)**.** Given a reconstruction functional $R : \mathcal{H} \to \mathcal{Q}$, two states $M, N \in \mathcal{H}$ are *memory-equivalent relative to* R if

$$M \sim_R N \iff d_{\mathcal{Q}}(R(M),\, R(N)) \leq \varepsilon.$$

A memory *persists* over $[t_0, t_1]$ when $M(t) \sim_R M(t_0)$ for all $t \in [t_0, t_1]$, regardless of whether the coefficients $a_i(t)$ or the material carriers of $M(t)$ change.

Proposition 7.3 (Memory Is Not Substrate Identity)**.** *If* $M(t) = \sum_i a_i(t)\psi_i$ *and* $N(t) = \sum_i b_i(t)\psi_i$ *satisfy* $d_{\mathcal{Q}}(R(M(t)), R(N(t))) \leq \varepsilon$*, then* $M(t)$ *and* $N(t)$ *instantiate the same recoverable memory relative to* R*, even if* $a_i(t) \neq b_i(t)$ *for every mode* i*.*

Proof. Memory equivalence (Definition 7.2) depends only on the image under R in $\mathcal{Q}$, not on equality in $\mathcal{H}$. Since $d_{\mathcal{Q}}(R(M(t)), R(N(t))) \leq \varepsilon$ holds by hypothesis, the two field states are memory-equivalent for every query stable under tolerance ε. Therefore memory persists as a recoverable distinction rather than as a conserved material or modal configuration. □

Chapter 8

Spectral Decomposition of Memory Fields

8.1 The Memory Operator

The modal expansion $M(t) = \sum_i a_i(t)\psi_i$ introduced in the previous chapter requires justification. The modes $\{\psi_i\}$ do not simply appear; they are eigenfunctions of a natural operator associated with the memory system. This chapter derives the modal basis from spectral theory, so that all subsequent modal machinery rests on a firm functional-analytic foundation.

Definition 8.1 (Memory Operator)**.** Let $\mathcal{H} = L^2(\Omega; \mathbb{C})$ be the memory field Hilbert space. A *memory operator* is a bounded, self-adjoint linear operator

$$\mathcal{L} : \mathcal{H} \longrightarrow \mathcal{H}$$

that encodes the intrinsic reconstructive structure of the cognitive domain Ω. Physically, $\mathcal{L}$ captures how the system propagates distinctions across space and frequency: high eigenvalues correspond to modes that are strongly supported by the substrate; low eigenvalues correspond to modes that are weakly supported and therefore easily overdamped.

Theorem 8.2 (Spectral Decomposition of $\mathcal{L}$)**.** *Let $\mathcal{L}$ be a compact, self-adjoint operator on $\mathcal{H}$. Then there exists a countable orthonormal family $\{\psi_i\}_{i\in\mathbb{N}} \subset \mathcal{H}$ and a sequence of real eigenvalues $\mu_1 \geq \mu_2 \geq \cdots \to 0$ such that*

$$\mathcal{L}\psi_i = \mu_i\,\psi_i \qquad \textit{for all } i \in \mathbb{N},$$

and every $M \in \mathcal{H}$ admits the expansion

$$M = \sum_{i=1}^{\infty} \langle M, \psi_i \rangle \, \psi_i,$$

convergent in the $\mathcal{H}$-norm.

Proof. By the spectral theorem for compact self-adjoint operators on a Hilbert space (see, e.g., Reed and Simon [20], Theorem VI.16), any compact self-adjoint operator admits a countable orthonormal eigenbasis with real eigenvalues accumulating only at zero. The expansion of M then follows from completeness of the eigenbasis. □

Definition 8.3 (Reconstructive Modes)**.** The eigenfunctions $\{\psi_i\}$ of $\mathcal{L}$ are the *reconstructive modes* of the memory system. The eigenvalue μ_i measures the *modal support*: the degree to which the substrate amplifies or sustains mode ψ_i. Modes with large μ_i are intrinsically stable; modes with small μ_i are weakly sustained and vulnerable to overdamping.

8.2 Time Evolution in the Modal Basis

Once the spectral basis is established, the dynamics of the memory field are naturally expressed modally.

Proposition 8.4 (Modal Coefficient Dynamics)**.** *Let $M(t) = \sum_i a_i(t)\psi_i$ be the modal expansion of the memory field, and suppose the field evolves by*

$$\frac{\partial M}{\partial t} = -\mathcal{L}^{\dagger} M + F(t),$$

where $\mathcal{L}^{\dagger}$ is a damping operator diagonal in the modal basis with entries $(\lambda_i + i\omega_i)$, and $F(t)$ is an external forcing term. Then each coefficient satisfies

$$\frac{da_i}{dt} = -(\lambda_i + i\omega_i)\, a_i + f_i(t),$$

where $f_i(t) = \langle F(t), \psi_i \rangle$.

Proof. Project the field equation onto mode ψ_i by taking the inner product with ψ_i:

$$\frac{d}{dt}\langle M(t), \psi_i \rangle = -\langle \mathcal{L}^{\dagger} M(t), \psi_i \rangle + \langle F(t), \psi_i \rangle.$$

Since $\mathcal{L}^\dagger$ is diagonal in the $\{\psi_i\}$ basis, $\langle \mathcal{L}^\dagger M, \psi_i \rangle = (\lambda_i + i\omega_i)\langle M, \psi_i \rangle = (\lambda_i + i\omega_i)a_i$. Substituting $a_i(t) = \langle M(t), \psi_i \rangle$ gives the stated ODE. □

This proposition retroactively justifies the "Damped Memory Mode" definition of Chapter 9: each coefficient ODE is not an assumption but a consequence of projecting the field equation onto the spectral basis supplied by Theorem 8.2.

8.3 The Modal Persistence Theorem

The spectral framework now enables a theorem that connects Part III directly to Part V. Persistence-through-flux (Chapter 13) shows that reconstructive density is conserved under balanced renewal. The following theorem shows that memory persists under arbitrary drift of modal coefficients, provided the reconstruction functional remains stable.

Theorem 8.5 (Modal Persistence)**.** *Suppose $M(t) = \sum_i a_i(t)\psi_i$ and let $R : \mathcal{H} \to \mathcal{Q}$ be a reconstruction functional that is continuous with respect to the $\mathcal{H}$-norm. If*

$$\|R(M(t)) - R(M(t_0))\|_{\mathcal{Q}} < \varepsilon$$

for all $t \in [t_0, t_1]$, then memory persists over $[t_0, t_1]$ despite arbitrary variation of the individual coefficients $a_i(t)$.

Proof. Memory persistence is defined by memory equivalence (Definition 7.2): $M(t) \sim_R M(t_0)$ whenever $d_{\mathcal{Q}}(R(M(t)), R(M(t_0))) \leq \varepsilon$. The hypothesis states exactly this condition. Therefore, regardless of how the coefficients $a_i(t)$ vary — they may oscillate, drift, grow, or shrink — the memory persists as long as the reconstruction functional's output remains within tolerance ε of its initial value. The coefficients are implementation details of the field; the reconstruction output is what constitutes the memory. □

Corollary 8.6 (Robustness of Memory to Substrate Fluctuation)**.** *Under the conditions of Theorem 8.5, a memory survives arbitrary fluctuation of the modal coefficients $a_i(t)$ provided the fluctuations collectively preserve the output of R. In particular,* *memory is robust to any transformation* $a_i \mapsto a_i + \delta_i$ *satisfying*

$$\left\| R\left(\sum_i (a_i + \delta_i)\psi_i\right) - R\left(\sum_i a_i\psi_i\right) \right\|_{\mathcal{Q}} < \varepsilon.$$

Proof. Substitute the perturbed coefficients into the condition of Theorem 8.5 and apply the definition of memory equivalence. □

The Modal Persistence Theorem is the bridge between the dynamical picture (modes oscillate and damp) and the ecological picture (memories survive through flux renewal): what matters is not the trajectory of any individual coefficient but the collective stability of what the reconstruction functional can recover.

Chapter 9

Dimensional Amplification and Witness Coordinates

9.1 Witness Expansion

Empirical work on the topological acoustic synapse (TAS) demonstrates that a small physical substrate driven by k input frequencies can generate an exponentially larger family of nonlinear interaction states — a *phi-bit lattice* — whose effective dimensionality far exceeds the input dimension [6]. This section formalizes the general principle as the *Dimensional Amplification Theorem.*

Definition 9.1 (Witness Expansion)**.** Let X be an input space of events, cues, or sensory fragments. A *witness expansion* is a map

$$\Phi : X \longrightarrow \mathbb{C}^n, \qquad x \mapsto (\phi_1(x), \dots, \phi_n(x)),$$

where each coordinate ϕ_j is a *reconstructive witness.* The expansion is *memory-useful* when there exists a low-complexity readout

$$L : \mathbb{C}^n \longrightarrow \mathcal{Q}$$

such that $L(\Phi(x)) \approx Q(x)$ for the relevant query $Q : X \to \mathcal{Q}$.

Definition 9.2 (Reconstructive Separability)**.** Two event classes $A, B \subset X$ are *reconstructively separable* by Φ when there exists a hyperplane vector $h \in \mathbb{C}^n$ such that

$$\mathrm{Re}\langle h, \Phi(a)\rangle > 0 \quad \text{and} \quad \mathrm{Re}\langle h, \Phi(b)\rangle < 0$$

for all $a \in A$ and $b \in B$ outside an admissible error set of measure zero.

Theorem 9.3 (Dimensional Amplification Theorem). *Let $\Phi : X \to \mathbb{C}^n$ be a nonlinear witness expansion. If two classes $A, B \subset X$ are not linearly separable in X but become linearly separable in $\Phi(X)$, then memory retrieval may be implemented by a simpler readout in witness space than in the original event space.*

Proof. Assume $A, B \subset X$ are not linearly separable in the original event coordinates. Suppose that under Φ there exists $h \in \mathbb{C}^n$ such that

$$\mathrm{sgnRe}\langle h, \Phi(x)\rangle$$

correctly distinguishes A from B. Then the decision boundary in X is the pullback

$$\Phi^{-1}\{z : \mathrm{Re}\langle h, z\rangle = 0\},$$

which may be a nonlinear hypersurface in X while remaining a hyperplane in $\Phi(X)$. The classification complexity has therefore been absorbed into the structure of the expansion Φ rather than the readout L. Retrieval is easier because the memory system searches the amplified witness space, not the original event space. □

Corollary 9.4 (Memory as Witness Redundancy). *If a memory is represented by n witnesses and reconstruction requires only $k < n$ witnesses, then the memory is robust to the loss of any $n - k$ witness coordinates, provided the remaining k coordinates determine the same query answer.*

Proof. Let $W = \{\phi_1, \ldots, \phi_n\}$ be the full witness set and $W' \subseteq W$ with $|W'| = k$. If $R_W(M) = R_{W'}(M)$ within tolerance ε, then deletion of the $n - k$ unused witnesses does not alter the recoverable distinction. Therefore the memory persists under partial witness loss whenever the remaining witnesses form a sufficient set for the query. □

Chapter 10

Resonant Modal Recall

10.1 Cue Coupling and Modal Dynamics

Functional MRI studies of intrinsic oscillatory modes reveal that large-scale brain activity decomposes as $\Psi(n,t) = \sum_\alpha \psi_\alpha(n)\tau_\alpha(t)$, where a small family of spatial modes ψ_α accounts for the majority of observed long-range connectivity [5]. MEM|8 reinterprets this finding: the modes are not incidental features of neural anatomy but the latent basis of a reconstructive memory field. Recall is resonance with that basis, not retrieval from a store.

Definition 10.1 (Cue Coupling)**.** A cue c *couples* to memory mode ψ_i with strength

$$\gamma_i(c) = |\langle C(c),\, \psi_i \rangle|,$$

where $C(c) \in \mathcal{H}$ is the field representation of the cue.

Definition 10.2 (Damped Memory Mode)**.** A memory mode has activation $a_i(t) \in \mathbb{C}$ governed by

$$\frac{da_i}{dt} = -(\lambda_i + i\omega_i)\, a_i + \gamma_i(c(t)) + \eta_i(t),$$

where $\lambda_i \geq 0$ is the *damping coefficient*, ω_i the *modal frequency*, $\gamma_i(c(t))$ is cue forcing, and $\eta_i(t)$ is background excitation (noise).

10.2 The Recall Threshold

Definition 10.3 (Ecphoric Threshold)**.** The *ecphoric threshold* $\theta_i > 0$ is the minimum activation magnitude required for mode ψ_i to contribute to conscious recall:

$$\text{mode } \psi_i \text{ is recall-active} \iff |a_i(t)| \geq \theta_i.$$

Theorem 10.4 (Resonant Recall Criterion)**.** *For constant cue forcing γ_i and negligible noise, the steady-state amplitude satisfies*

$$|a_i^*| = \frac{|\gamma_i|}{\sqrt{\lambda_i^2 + \omega_i^2}}.$$

Recall occurs precisely when

$$|\gamma_i| \;\geq\; \theta_i \sqrt{\lambda_i^2 + \omega_i^2}.$$

Proof. For constant forcing the ODE reduces to

$$\frac{da_i}{dt} = -(\lambda_i + i\omega_i)\, a_i + \gamma_i.$$

The unique equilibrium is

$$a_i^* = \frac{\gamma_i}{\lambda_i + i\omega_i}.$$

Taking magnitudes:

$$|a_i^*| = \frac{|\gamma_i|}{|\lambda_i + i\omega_i|} = \frac{|\gamma_i|}{\sqrt{\lambda_i^2 + \omega_i^2}}.$$

Recall requires $|a_i^*| \geq \theta_i$, which gives

$$|\gamma_i| \geq \theta_i \sqrt{\lambda_i^2 + \omega_i^2},$$

establishing the threshold condition. □

Remark 10.5. Recall is not lookup. It is thresholded resonant excitation. Two factors govern whether a cue triggers recall: cue-mode overlap (the numerator $|\gamma_i|$) and modal impedance (the denominator $\sqrt{\lambda_i^2 + \omega_i^2}$). Strong cues with high coupling always recall; weak or mismatched cues fail, even when the mode itself is intact.

Chapter 11

Forgetting as Overdamping, Repair as Resonance Restoration

11.1 Overdamped Modes

Definition 11.1 (Overdamped Memory)**.** A memory mode ψ_i is *overdamped relative to a cue class* $\mathcal{C}$ when

$$\sup_{c \in \mathcal{C}} \frac{|\gamma_i(c)|}{\sqrt{\lambda_i^2 + \omega_i^2}} < \theta_i.$$

Proposition 11.2 (Forgetting Is Loss of Excitability)**.** *If a mode is overdamped relative to all admissible cues, the corresponding memory is operationally forgotten but structurally latent.*

Proof. The mode ψ_i remains an element of the modal basis $\{\psi_i\}$; it is not removed from $\mathcal{H}$. However, for every $c \in \mathcal{C}$,

$$|a_i^*(c)| < \theta_i,$$

so the mode cannot become recall-active. Forgetting is therefore a failure of excitability, not an erasure of the structural possibility of the memory. □

This distinction is philosophically important. A forgotten memory is not destroyed; it is unreachable. The distinction between erasure and overdamping generates the possibility of recovery through repair.

11.2 Repair as Damping Reduction

Theorem 11.3 (Repair by Reducing Damping)**.** *Let ψ_i be overdamped under forcing γ_i. If repair changes the damping coefficient from λ_i to λ_i' with*

$$\lambda_i' < \sqrt{\left(\frac{|\gamma_i|}{\theta_i}\right)^2 - \omega_i^2},$$

then the memory becomes recall-active under the same cue forcing.

Proof. The repaired steady-state amplitude is

$$|a_i^{*\prime}| = \frac{|\gamma_i|}{\sqrt{(\lambda_i')^2 + \omega_i^2}}.$$

The stated inequality on λ_i' is equivalent to

$$(\lambda_i')^2 + \omega_i^2 < \left(\frac{|\gamma_i|}{\theta_i}\right)^2.$$

Taking square roots and rearranging:

$$\frac{|\gamma_i|}{\sqrt{(\lambda_i')^2 + \omega_i^2}} > \theta_i.$$

Therefore $|a_i^{*\prime}| > \theta_i$, meaning the repaired mode is recall-active. □

Corollary 11.4. *Therapeutic, environmental, and social interventions that restore memory do not reconstruct stored content. They reduce the effective damping coefficient of overdamped modes, allowing existing latent structure to become resonant.*

Part IV

Ecological Memory

Chapter 12

Memory Outside the Brain

12.1 The Ecological Extension

The neural interpretation of MEM|8 is a special case, not the general case. Any system whose present state is a compressed witness of its event history satisfies the axioms of the MEM|8 framework. Memory is therefore not confined to neural tissue.

Definition 12.1 (Environmental Memory Field)**.** An *environmental memory field* is a spatial distribution

$$\Phi_M : \Omega \longrightarrow \mathbb{R}_{\geq 0}$$

representing the *local reconstructive support* at each point $x \in \Omega$. High values of $\Phi_M(x)$ indicate regions where distinction reconstruction is strongly supported; low values indicate reconstructive poverty.

Example 12.2 (Workspace as Memory Architecture)**.** Consider a knowledge worker's physical desk. Let Ω be the desk surface. Notes, open books, pinned diagrams, and unfinished drafts each instantiate a locally high Φ_M. The workspace as a whole is a distributed memory architecture in which cognitive load is partially offloaded into the environmental field.

Proposition 12.3 (Ecological Extension of the Memory Field)**.** *Let a cognitive domain extend across an agent A and an environment* Env*. The composite memory field is*

$$M_{total}(t) = M_A(t) + M_{\mathrm{Env}}(t),$$

and the composite modal decomposition is

$$M_{total}(t) = \sum_i a_i^A(t)\psi_i^A + \sum_j a_j^{\mathrm{Env}}(t)\psi_j^{\mathrm{Env}}.$$

Recall can be triggered by cue coupling to either component.

Proof. Both components are elements of the same Hilbert space $\mathcal{H}$ (extended to the composite domain). Their sum is again an element of $\mathcal{H}$, with coefficients summing over the two mode families. Recall via Theorem 10.4 applies to any mode in the combined basis. □

Chapter 13

Habitat Versus Archive

Definition 13.1 (Archive)**.** An *archive* is a system designed to preserve the information content of representations across time. Its memory metric is storage fidelity: $d_{\mathcal{Q}}(R_t(S(t)), q_0)$ should remain small as t increases.

Definition 13.2 (Habitat)**.** A *habitat* is a system designed to preserve and support the *reconstruction pathways* through which distinctions can be regenerated. Its memory metric is reconstructive accessibility: the probability that a randomly sampled cue activates a recall-active mode.

Proposition 13.3 (Archives Can Fail as Habitats)**.** *An archive may preserve complete information while failing as a memory habitat if the reconstruction pathways required to use that information are lost.*

Proof. Consider a text perfectly preserved in a dead language. Every symbol is retained; $d_{\mathcal{Q}}$ of the stored content is zero. Yet if no agent can execute the reconstruction operator R_t that maps the symbols to their distinctions, the memory is operationally inaccessible. The archive is intact; the habitat is broken. □

This distinction has consequences for the design of libraries, databases, scientific journals, and institutional knowledge systems. Preservation of content is necessary but insufficient. Preservation of the reconstruction community, the interpretive tradition, and the cue structures that trigger recall is equally essential.

Part V

Memory and Repair

Chapter 14

Memory as Flux Maintenance

14.1 Persistence Through Replacement

Synaptic proteins mediating vesicle exocytosis have average lifetimes of one to two days, yet the memories they support persist for decades [18]. This is the biological instance of the ship-of-Theseus problem: how can identity persist when every component is replaced?

MEM|8 resolves the paradox by locating memory identity in reconstructive organization rather than material constitution. The mathematical expression is a flux-balance equation.

Definition 14.1 (Memory Flux Balance)**.** Let $\rho(x,t)$ be the local *reconstructive density* at point $x \in \Omega$. Let $J(x,t)$ be the flux of reconstructive carriers, S_{renew} a renewal source term, and S_{decay} a degradation sink term. The *memory balance equation* is

$$\frac{\partial \rho}{\partial t} + \nabla \cdot J = S_{\text{renew}} - S_{\text{decay}}.$$

Theorem 14.2 (Persistence Through Replacement)**.** *A memory region $U \subset \Omega$ preserves its reconstructive density over time if*

$$\int_U (S_{\text{renew}} - S_{\text{decay}})\, dx = \int_{\partial U} J \cdot \hat{n}\, dS,$$

in which case $\dfrac{d}{dt}\displaystyle\int_U \rho\, dx = 0.$

Proof. Integrate the balance equation over U:

$$\frac{d}{dt}\int_U \rho\, dx = -\int_{\partial U} J \cdot \hat{n}\, dS + \int_U (S_{\text{renew}} - S_{\text{decay}})\, dx.$$

By the stated equality, the two terms on the right cancel, giving $\frac{d}{dt}\int_U \rho\, dx = 0$. $\square$

Corollary 14.3 (Ship-of-Theseus Memory). *If every material carrier in U is replaced over the interval $[0, T]$ but the flux-balance condition holds throughout, then the memory persists as a stable reconstructive organization.*

Proof. Theorem 14.2 guarantees that total reconstructive density in U is conserved. Memory identity is defined by recoverability (Definition 7.2), not by carrier identity. Therefore wholesale replacement of carriers is compatible with persistent memory. $\square$

14.2 Two Transport Modes: Circulation and Capture

Definition 14.4 (Two-State Memory Transport). Memory carriers occupy either a *circulating state* C (preserving global availability) or a *captured state* K (stabilizing local reconstruction). Let $p_K(v, d)$ be the probability that a carrier of speed v moving in direction $d \in \{+, -\}$ is captured by a local memory site.

Assumption 14.5 (Speed-Dependent Capture). There exists a threshold v_c such that

$$v < v_c \;\Rightarrow\; p_K(v, d) \approx p_{\text{slow}}, \qquad v > v_c \;\Rightarrow\; p_K(v, d) \approx p_{\text{fast}},$$

with $p_{\text{slow}} > p_{\text{fast}}$.

Theorem 14.6 (Capture-Stabilization Principle). *Under Assumption 14.5, slow carriers preferentially stabilize local memory sites while fast carriers preferentially maintain global circulation.*

Proof. Let $N = N_{\text{slow}} + N_{\text{fast}}$. Expected capture counts are

$$\mathbb{E}[K] = p_{\text{slow}} N_{\text{slow}} + p_{\text{fast}} N_{\text{fast}},$$

$$\mathbb{E}[C] = (1 - p_{\text{slow}}) N_{\text{slow}} + (1 - p_{\text{fast}}) N_{\text{fast}}.$$

Since $p_{\text{slow}} > p_{\text{fast}}$, slow carriers contribute disproportionately to $\mathbb{E}[K]$ and fast carriers to $\mathbb{E}[C]$. The two velocity classes therefore separate into local stabilization and global circulation roles. $\square$

Theorem 14.7 (Retrograde Renewal Bias)**.** *Suppose carriers moving retrograde have lower capture probability than anterograde carriers: $p_K(v,-) < p_K(v,+)$. Then, even when initial directional populations are equal, the surviving traversing flux has a retrograde bias.*

Proof. Assume $N_+ = N_-$. Expected traversing populations are

$$T_+ = (1 - p_K(v,+))N_+, \qquad T_- = (1 - p_K(v,-))N_-.$$

Since $p_K(v,-) < p_K(v,+)$ implies $1 - p_K(v,-) > 1 - p_K(v,+)$, and $N_+ = N_-$, we obtain $T_- > T_+$. Retrograde bias in the traversing flux arises from differential capture probability, not from directional motor asymmetry. □

Theorem 14.7 gives MEM|8 a mechanism for automatic renewal: old material is preferentially returned for replacement without requiring explicit deletion or directed degradation signals.

Chapter 15

Forgetting as Structural Deformation

Definition 15.1 (Deformation Cost)**.** The *deformation cost* of the memory field over the interval $[t, t + \Delta]$ is

$$D(H_t, H_{t+\Delta}) = \sum_{i \in I} |a_i(t + \Delta) - a_i(t)|^2 \cdot \mathbf{1}[|a_i(t)| \geq \theta_i].$$

The indicator selects only modes that were recall-active at time t: only the deformation of active memories contributes.

Proposition 15.2 (Forgetting Rate Bounds Deformation Cost)**.** *Let* $\Lambda = \sup_i \sqrt{\lambda_i^2 + \omega_i^2}$ *be the maximum modal impedance over all recall-active modes, and suppose no cue forcing is present. Then*

$$D(H_t, H_{t+\Delta}) \leq \Lambda^2 \cdot \Delta^2 \cdot \|M(t)\|^2_{\text{active}},$$

where $\|M(t)\|^2_{\text{active}} = \sum_{i\colon |a_i| \geq \theta_i} |a_i(t)|^2$.

Proof. Without cue forcing and noise, $a_i(t) = a_i(0)e^{-(\lambda_i + i\omega_i)t}$. Over a short interval Δ, differentiating gives $\dot{a}_i = -(\lambda_i + i\omega_i)a_i$, so

$$|a_i(t + \Delta) - a_i(t)| \approx |\dot{a}_i(t)|\, \Delta = |\lambda_i + i\omega_i|\, |a_i(t)|\, \Delta = \sqrt{\lambda_i^2 + \omega_i^2}\, |a_i(t)|\, \Delta \leq \Lambda\, |a_i(t)|\, \Delta,$$

using the definition $\Lambda = \sup_i \sqrt{\lambda_i^2 + \omega_i^2}$. Squaring and summing over recall-active modes:

$$D(H_t, H_{t+\Delta}) = \sum_{i\colon |a_i| \geq \theta_i} |a_i(t + \Delta) - a_i(t)|^2 \leq \Lambda^2 \Delta^2 \sum_{i\colon |a_i| \geq \theta_i} |a_i(t)|^2 = \Lambda^2 \Delta^2 \|M(t)\|^2_{\text{active}}.$$

□

Chapter 16

Persistent Anomalies and Scientific Memory

Definition 16.1 (Anomaly Persistence)**.** An anomaly $\mathcal{A}$ in a knowledge system is *persistent* if its associated memory mode $\psi_{\mathcal{A}}$ has low damping coefficient $\lambda_{\mathcal{A}} \approx 0$ and strong coupling to routine cues: $\gamma_{\mathcal{A}}(c) \gg \theta_{\mathcal{A}}\sqrt{\lambda_{\mathcal{A}}^2 + \omega_{\mathcal{A}}^2}$ for most c in the standard repertoire.

Proposition 16.2 (Anomaly Survives Paradigm Shift)**.** *If an anomaly has low damping and strong cue coupling, it remains recall-active across theoretical transitions that alter other modes.*

Proof. Theoretical transitions change the modal basis $\{\psi_i\} \to \{\psi_i'\}$ and alter the damping landscape for most modes. But if $\lambda_{\mathcal{A}} \approx 0$, the anomaly mode remains below the recall threshold even under strong perturbation of the remaining basis. Its cue coupling $\gamma_{\mathcal{A}}$ ensures routine re-excitation. Therefore the anomaly persists: it is the memory mode most resistant to paradigm-induced overdamping. $\square$

Science accumulates not because it stores truths, but because the anomaly modes of the scientific community are collectively maintained at low damping through continual cue excitation: publication, teaching, replication, and debate.

Part VI

Computation and MEM|8

Chapter 17

Memory Machines

Definition 17.1 (Reconstruction Capacity)**.** The *reconstruction capacity* of a computational memory system is the maximum number of distinct distinctions recoverable from the system's current state, measured in bits:

$$\mathcal{RC}(S) = \log_2 |\{q \in \mathcal{Q} : \exists R,\, R(S) = q\}|.$$

Remark 17.2. Classical RAM measures memory in storage capacity (bits of information stored). MEM|8 measures memory in reconstruction capacity (bits of information recoverable). The two coincide for perfect-fidelity storage; they diverge for systems with compression, degradation, and reconstruction.

Theorem 17.3 (Reconstruction Dominates Storage)**.** *For a system implementing witness expansion $\Phi : X \to \mathbb{C}^n$ with n independent binary witnesses, the reconstruction capacity satisfies*

$$\mathcal{RC}(S) \ \geq \ n,$$

independently of the physical storage capacity of the input representation.

Proof. Each witness coordinate $\phi_j : X \to \mathbb{C}$ contributes at least one bit of distinguishing power: whether $\mathrm{Re}(\phi_j(x)) > 0$ or not partitions X into two classes. With n independent witnesses, the number of jointly distinguishable classes is at least 2^n. The reconstruction capacity is defined as $\mathcal{RC}(S) = \log_2 |\{q \in \mathcal{Q} : \exists R,\, R(S) = q\}|$, so having at least 2^n distinguishable outputs gives

$$\mathcal{RC}(S) \ \geq \ \log_2 2^n = n.$$

This lower bound is determined by the expansion dimension n, not the input dimension $\dim(X)$. $\square$

Chapter 18

Event-Log Computation and Spherepop

Programs in the event-log interpretation are not state machines; they are history machines. A program execution is a sequence of events $E = (e_1, \ldots, e_n)$ and the result is a recovered distinction from the terminal historical state H_n.

This connects directly to the Spherepop calculus, in which computation proceeds through irreversible events (Pop, Refuse, Collapse, Bind) and the computational state is the history of those events rather than a mutable store.

Definition 18.1 (Event-Log Program)**.** An *event-log program* is a pair $(\mathcal{E}, R)$ where $\mathcal{E}$ is an event alphabet and $R : \mathcal{H}_{\mathcal{E}} \to \mathcal{Q}$ is a reconstruction functional over histories in $\mathcal{H}_{\mathcal{E}}$. Execution consists of generating a history H; the result is $R(H)$.

Proposition 18.2 (Spherepop as Event-Log Computation)**.** *The Spherepop calculus, with its irreversible event operators, is an event-log program in the sense above, where the history H is the sequence of sphere-state transitions and R maps terminal histories to denotational values.*

Sketch. Each Spherepop operator (Pop, Refuse, Collapse, Bind) modifies the global sphere configuration irreversibly, generating a new event in the history. The operational semantics assigns meaning to terminal histories. This is precisely the structure of an event-log program. □

Chapter 19

A Concrete Instantiation: The MEM|8 Rust Backend

19.1 From Theory to Implementation

The preceding chapters developed MEM|8 as an abstract mathematical framework: memory fields in $L^2(\Omega;\mathbb{C})$, modal expansions, resonant recall, and flux-balanced persistence. This chapter examines a working Rust implementation of MEM|8's storage layer [24] and shows precisely where the abstract constructions land in code. The implementation is instructive not only as an existence proof but because it makes visible one residual tension: the `VolatileBackend`/`PersistentBackend` split still carries traces of the storage illusion that MEM|8 theoretically supersedes.

19.2 Wave Coordinates as Witness Expansion

The core type is

```
pub type WaveCoord = (u8, u8, u16);
```

a triple (x, y, f) where $x, y \in \{0, \ldots, 255\}$ specify a position in the 256×256 spatial grid and $f \in \{0, \ldots, 65535\}$ specifies a frequency band. The full coordinate space has cardinality

$$256 \times 256 \times 65536 = 2^{32} \approx 4.3 \times 10^9$$

distinct positions, each capable of holding a complex amplitude

```
pub type Wave = Complex64;
```

This is a concrete realization of the witness expansion (Definition 9.1). Each coordinate $\phi_{(x,y,f)}$ is a reconstructive witness, and the full lattice

$$\Lambda = \{(x, y, f) : x, y \in [0, 255],\ f \in [0, 65535]\}$$

is the high-dimensional witness space into which events are expanded. Theorem 9.3 applies directly: input events that are not linearly separable in raw sensory space may become separable in Λ under a nonlinear embedding, and Corollary 9.4 guarantees robustness to the loss of individual witness coordinates.

The choice of `u8` for spatial axes and `u16` for frequency is itself a statement about the relative grain of spatial versus spectral resolution in the intended cognitive domain, an architectural commitment to the *type-specific encoding* principle: different kinds of distinctions (spatial vs. spectral) occupy different regions of the witness lattice.

19.3 `query_region` as Ecphoric Activation

The central method of the `WaveStorageBackend` trait is

```
fn query_region(&self, bounds: BoundingBox)
    -> Result<Vec<(WaveCoord, Wave)>>;
```

where `BoundingBox` specifies spatial extent:

```
pub struct BoundingBox {
    pub min_x: u8, pub min_y: u8,
    pub max_x: u8, pub max_y: u8,
}
```

The implementation comment reads: "This is the foundation of resonance. The backend returns all waves within the spatial bounds, which the cognitive layer then uses to calculate interference patterns."

This maps precisely onto the ecphoric activation operator of Chapter 7. A cue c is encoded as a bounding box $B_c \subset \Lambda$. The cue coupling of c to mode ψ_i is

$$\gamma_i(c) = \big|\langle C(c),\, \psi_i \rangle\big|,$$

where $C(c)$ is the field representation of the bounding-box query: the indicator function of B_c in the wave lattice. The backend computes, for each coordinate $(x, y, f) \in B_c$, the stored complex amplitude $a_{(x,y,f)}$. These are exactly the excitation coefficients $a_i(t)$ of Definition 7.1, restricted to the modes whose spatial support intersects the cue region.

Resonant recall (Theorem 10.4) then operates in the cognitive layer above the backend. The cognitive core receives the amplitudes returned by `query_region`, computes interference patterns, and applies a threshold: a mode contributes to recall if and only if

$$|a_i^*| \;\geq\; \theta_i.$$

The separation of concerns is therefore: the backend handles wave retrieval (which modes are active in the cue region); the cognitive layer handles thresholding and interference (which of those modes cross the recall boundary). This is the correct factoring — the storage substrate should not know the ecphoric threshold, just as a physical medium should not know which patterns are meaningful to the organism reading from it.

Proposition 19.1 (`query_region` Instantiates Ecphory)**.** *Let B_c be a bounding box encoding cue c, and let $\mathcal{W}(B_c)$ denote the set of wave amplitudes returned by* `query_region`(B_c)*. Then recall of memory m from cue c is equivalent to*

$$\left| \sum_{(x,y,f)\,\in\,B_c} a_{(x,y,f)}\,\psi_{(x,y,f)} \right| \;\geq\; \theta_m,$$

where $\psi_{(x,y,f)}$ are the lattice basis modes and θ_m is the ecphoric threshold for m.

Proof. `query_region` returns exactly the set $\{(\mathit{coord},\, a_{\mathit{coord}})\}$ for $\mathit{coord} \in B_c \cap \Lambda$. The cognitive layer forms the superposition

$$\hat{M}(c) = \sum_{\mathit{coord}\,\in\,B_c} a_{\mathit{coord}}\,\psi_{\mathit{coord}},$$

which is the restriction of the memory field $M(t)$ to the cue region — the ecphoric activation $\mathcal{E}(c, H) \to \hat{H}$ of Chapter 7, with the bounding box selecting the relevant historical submanifold. Recall fires when $|\hat{M}(c)| \geq \theta_m$, matching Theorem 10.4. $\square$

19.4 The Storage Illusion Residue

The `flush()` method and the deprecated `WaveStorage` trait reveal a residual tension between the implementation and the theory.

`flush()` is described as ensuring durability by flushing write buffers to persistent storage, and as a no-op for volatile backends. This treats persistence as a property of the backend type (volatile vs. persistent) rather than as a property of the maintenance dynamics.

The MEM|8 theory says something different. Persistence is not a function of where data sits (RAM versus disk) but of whether the flux-balance condition of Theorem 14.2 is satisfied:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot J = S_{\text{renew}} - S_{\text{decay}}.$$

A "persistent" backend that is never re-excited by cue queries will accumulate overdamped modes and lose memory in the operational sense, even though every byte is preserved. Conversely, a "volatile" backend continually re-queried and refreshed is maintaining its reconstructive density through active use — a form of the renewal flux S_{renew} demonstrated in the synaptic-vesicle trafficking model of Chapter 13.

The deprecated `WaveStorage` / `WaveGrid` pair is a clean example of the storage illusion at the implementation level. `store_wave_grid` writes a static grid and returns; there is no provision for re-excitation, ecphoric thresholds, or flux maintenance. Its deprecation in favor of `WaveStorageBackend` with `query_region` at the center is, in miniature, the same conceptual transition that MEM|8 makes at the theoretical level: from storage as the primitive operation to reconstruction as the primitive operation.

19.5 Toward a Fully Theory-Conformant Backend

A backend fully aligned with MEM|8 theory would augment `WaveStorageBackend` with three additional mechanisms.

Damping and renewal. A method `decay_step(&mut self, dt: f64)` would apply

$$a_i(t + \Delta t) = a_i(t)\, e^{-(\lambda_i + i\omega_i)\Delta t}$$

to every stored coefficient, and a companion `renew(&mut self, flux: &[(WaveCoord, Wave)])` would add renewal input S_{renew} to specified coordinates, making persistence a dynamic equilibrium rather than a backend attribute.

Damping coefficient registry. Each coordinate (x, y, f) should carry $\lambda_{(x,y,f)}$ and $\omega_{(x,y,f)}$, so the cognitive layer can evaluate the recall threshold condition

$$|\gamma_i| \;\geq\; \theta_i \sqrt{\lambda_i^2 + \omega_i^2}$$

without embedding modal parameters in the storage layer.

Retrograde sweep. A scheduled low-priority process implementing Theorem 14.7 would flag low-activity coordinates for renewal before they become overdamped, maintaining the maintenance current that keeps long-term memories recoverable without explicit refresh commands.

These three additions close the gap between the current implementation and the theoretical architecture. The implementation already has the right shape at the `query_region` level; what remains is to make the temporal dynamics — damping, renewal, retrograde sweep — first-class citizens of the backend interface rather than implicit assumptions about the environment.

Chapter 20

The Halting Problem and Historical Persistence

Theorem 20.1 (Undecidability as Future Reconstructability Failure)**.** *The halting problem expresses a fundamental limit on future reconstructability: there is no finite procedure that determines, for an arbitrary event-log program and initial condition, whether the history will eventually produce a terminal reconstruction.*

Reduction to Rice's Theorem. Any property of the language recognized by a Turing machine (including whether the machine halts on a given input) is a property of the reconstruction functional $R : \mathcal{H} \to \mathcal{Q}$ over all possible histories. By Rice's theorem, every nontrivial such property is undecidable. Therefore the question "will the reconstruction eventually succeed?" is undecidable for arbitrary programs. $\square$

The interpretation is important: undecidability is not a failure of storage but a failure of *future reconstructability.* We cannot predict whether the event-log will eventually support the desired reconstruction. This gives the halting problem a MEM|8 reading entirely consistent with the general framework.

Part VII

Society and Civilization

Chapter 21

Institutional Memory

Definition 21.1 (Organizational Memory Field)**.** The *organizational memory field* of an institution is the composite memory field

$$M_{\mathrm{org}}(t) = \sum_{a \in A} M_a(t) + M_{\mathrm{doc}}(t) + M_{\mathrm{proc}}(t),$$

where the sum is over individual agents $a \in A$, plus document repositories M_{doc} and procedural structures M_{proc}.

Proposition 21.2 (Personnel Change Does Not Destroy Memory)**.** *A personnel change (departure of agent a) destroys organizational memory if and only if the modes $\{\psi_i^a\}$ carried by agent a are not decodable from the remaining composite field $M_{\mathrm{org}} - M_a$.*

Proof. By Corollary 9.4, a memory survives the loss of witness coordinates provided sufficient redundant witnesses remain. If M_{doc} or M_{proc} supply alternative cue paths to the same modal distinctions, the reconstruction persists. Memory loss occurs precisely when no such redundant pathway exists. □

Chapter 22

Scientific Memory

Science is a civilization-scale memory architecture. Its function is to maintain, across generations and institutional boundaries, a reconstructive field rich enough to support cumulative inquiry.

Definition 22.1 (Scientific Paper as Reconstruction Operator)**.** A scientific paper P is a reconstruction operator $R_P : \mathcal{H}_{\text{reader}} \to \mathcal{Q}_{\text{science}}$ that maps the reader's current cognitive state to a set of distinctions about the natural world.

Proposition 22.2 (Citation Network as Modal Coupling)**.** *The citation network of the scientific literature implements a system of modal couplings γ_{ij} between papers P_i and P_j: a citation from P_j to P_i provides a coupling channel that keeps the modes of P_i excitable via engagement with P_j.*

Proof. Reading P_j generates cue coupling to the modes encoded in P_i through the citation-link (references, shared terminology, shared problems). This is cue forcing γ_i in the sense of Definition 7.2, applied to the scientific memory field. $\square$

Chapter 23

Cultural Memory

Languages, traditions, norms, and rituals are long-lived repair structures. They persist by maintaining the reconstructive conditions under which the distinctions they encode can be regenerated.

Theorem 23.1 (Cultural Memory Survival Condition)**.** *A cultural form $\mathcal{F}$ survives across generations if and only if the flux-balance condition of Theorem 14.2 is maintained for the reconstructive density of $\mathcal{F}$: renewal (teaching, practice, performance, reenactment) balances decay (mortality, forgetting, disuse).*

Proof. Cultural forms are event-log systems (Chapter 6). Their memory field satisfies the balance equation (Definition 14.1). Survival of the form over time corresponds to $\frac{d}{dt}\int \rho\, dx = 0$, which by Theorem 14.2 holds if and only if the balance condition holds. $\square$

Part VIII

Foundations

Chapter 24

Memory Before Representation

The standard order of explanation places representation first. There are things; representations encode them; memory stores the encodings.

MEM|8 reverses this order at every level.

Theorem 24.1 (Representation Is Derived)**.** *A representation r of a distinction δ exists precisely when there is a reconstruction operator R and a substrate S such that $R(S) = \delta$. Representations are not primitive; they are the outputs of reconstruction.*

Proof. By Definition 1.4, reconstructive capacity is defined without appeal to representations. A representation can then be defined as a naming of the output of a reconstruction operator: $r \equiv R(S)$. This makes representation derivable from recoverability, not the reverse. $\square$

The inversion is therefore:

$$\text{Traditional:} \quad \text{Representation} \to \text{Storage} \to \text{Memory}$$

$$\text{MEM|8:} \quad \text{Events} \to \text{Recoverability} \to \text{Memory} \to \text{Representation}$$

Chapter 25

Memory Before Identity

Definition 25.1 (Identity as Stable Reconstruction)**.** The *identity* of a system Σ over an interval $[t_0, t_1]$ is the equivalence class of states under memory-equivalence relative to a canonical reconstruction operator R_Σ:

$$\mathrm{Id}(\Sigma, [t_0, t_1]) = [S(t_0)]_{\sim_{R_\Sigma}}.$$

Identity persists over the interval if $S(t) \sim_{R_\Sigma} S(t_0)$ for all $t \in [t_0, t_1]$.

Proposition 25.2 (The Self as Persistent Memory Process)**.** *Personal identity over a lifetime is the persistence of a memory field under continual flux, satisfying the balance condition of Theorem 14.2.*

Proof. Personal identity is a special case of system identity. The canonical reconstruction operator R_{self} maps current neural and embodied state to autobiographical distinctions. Persistence requires only that the flux-balance condition hold for the reconstructive density of the autobiographical memory field — not that any neural substrate remain unchanged. This is consistent with the empirical observation that virtually every molecule in the human body is replaced over a decade, while autobiographical memory persists. $\square$

Chapter 26

Memory Before Prediction

This chapter addresses an implication the preceding chapters repeatedly approach but never state outright. If memory is recoverable continuation, then prediction — the anticipation of future distinctions — is a special case of reconstruction applied forward rather than backward. Prediction presupposes memory, not the reverse.

26.1 Prediction as Forward Reconstruction

Definition 26.1 (Predictive Reconstruction)**.** A *predictive reconstruction* is a reconstruction operator R_+ that maps the current memory field $M(t)$ to an anticipated future distinction:

$$R_+ : \mathcal{H} \longrightarrow \mathcal{Q}, \qquad M(t) \mapsto \hat{q}(t+\tau),$$

where $\tau > 0$ is the prediction horizon and $\hat{q}(t+\tau)$ is the anticipated state of the distinction space at time $t+\tau$.

Theorem 26.2 (Prediction Presupposes Memory)**.** *A system capable of predictive reconstruction R_+ must have nonzero reconstructive capacity $\mathcal{RC}(S) > 0$.*

Proof. Predictive reconstruction R_+ maps $M(t)$ to a future distinction $\hat{q}(t+\tau)$. For R_+ to be non-trivial (i.e., better than chance), it must discriminate among distinct current states: if $\mathcal{RC}(S) = 0$, every state maps to the same reconstruction output (Theorem 3.2), so no state-dependent prediction is possible. Therefore non-trivial prediction requires $\mathcal{RC}(S) > 0$, i.e., genuine reconstructive capacity over the current memory field. Prediction is downstream of memory, not independent of it. □

Corollary 26.3 (No Prediction Without Historical State). *A memoryless system — one with $H_t = \varnothing$ for all t — cannot predict.*

Proof. Without historical state there is no compression to unpack (Chapter 5), no modal basis to excite (Chapter 9), and no reconstruction functional to apply. Theorem 26.2 then gives $\mathcal{RC}(S) = 0$, so no non-trivial prediction is possible. □

26.2 Predictive Coding as Memory Resonance

The predictive processing framework (Friston [11]) treats perception and action as continuous minimization of prediction error: the brain maintains a generative model and updates it when predictions fail.

Under MEM|8 this has a natural modal reading. The generative model is the current modal excitation pattern $\{a_i(t)\}$. A prediction is a projection of that pattern forward:

$$\hat{a}_i(t+\tau) = a_i(t)\, e^{-(\lambda_i + i\omega_i)\tau}.$$

Prediction error is the discrepancy between the predicted field $\hat{M}(t+\tau) = \sum_i \hat{a}_i(t+\tau)\psi_i$ and the actual field $M(t+\tau)$:

$$\mathcal{E}_{\text{pred}}(t,\tau) = \|M(t+\tau) - \hat{M}(t+\tau)\|^2_{\mathcal{H}} = \sum_i |a_i(t+\tau) - \hat{a}_i(t+\tau)|^2.$$

Proposition 26.4 (Prediction Error Bounds Surprise). *Let $\mathcal{E}_{\text{pred}}(t,\tau)$ be the prediction error defined above. Then*

$$\mathcal{E}_{\text{pred}}(t,\tau) \le \sum_i |\eta_i(t,\tau)|^2 + \sum_i |\gamma_i(c)|^2,$$

where $\eta_i(t,\tau)$ is background noise accumulated over $[t, t+\tau]$ and $\gamma_i(c)$ is the cue forcing from unexpected environmental input over the same interval.

Proof. The actual coefficient satisfies $a_i(t+\tau) = \hat{a}_i(t+\tau) + \delta_i$, where δ_i arises from noise η_i and unexpected forcing $\gamma_i(c)$ not already encoded in the current modal state. Therefore $|a_i(t+\tau) - \hat{a}_i(t+\tau)|^2 = |\delta_i|^2 \le |\eta_i|^2 + |\gamma_i(c)|^2$, and summing over modes gives the bound. □

Prediction error is thus bounded by noise plus unencoded surprise. A system with richer memory (higher-dimensional modal basis, lower damping) encodes more of its environment and therefore anticipates more accurately — connecting MEM|8 directly to the Bayesian brain hypothesis and active inference.

26.3 Reservoir Computing as MEM|8 Instantiation

Reservoir computing [?] treats a high-dimensional dynamical system (the reservoir) as a fixed feature expander: inputs drive the reservoir, and a simple linear readout is trained on the reservoir state to perform prediction and classification tasks.

This is precisely the MEM|8 architecture of Theorem 9.3:

$$\begin{aligned}
\text{Reservoir state} &\longleftrightarrow \text{Witness expansion } \Phi(x) \in \mathbb{C}^n \\
\text{Linear readout} &\longleftrightarrow \text{Reconstruction functional } L : \mathbb{C}^n \to \mathcal{Q} \\
\text{Prediction task} &\longleftrightarrow \text{Predictive reconstruction } R_+
\end{aligned}$$

The reservoir's memory of past inputs (its "echo state") is the historical state H_t; the readout extracts distinctions from that compressed history. Theorem 26.2 explains why reservoir memory capacity directly governs prediction accuracy: without sufficient $\mathcal{RC}(S)$, the readout cannot form useful predictions.

The TAS phi-bit lattice of Chapter 8 is a physical instantiation of a reservoir operating in the MEM|8 regime: a small substrate generates a large witness space, and classification (prediction) becomes linearly separable in witness coordinates.

26.4 The Inversion

The standard order places prediction first: a system makes predictions; successful prediction constitutes understanding; memory is the residue of past understanding.

MEM|8 inverts this:

$$\text{Traditional:} \quad \text{Prediction} \to \text{Model} \to \text{Memory}$$

$$\text{MEM|8:} \quad \text{Memory} \to \text{Modal Structure} \to \text{Prediction}$$

Prediction is reconstruction applied in the forward temporal direction. It inherits all the properties of reconstruction: it is approximate rather than exact, threshold-gated rather than continuous, and ecologically extended across agent and environment rather than confined to a brain.

The arrow of prediction points forward in time; the arrow of memory points backward; both are instantiated by the same reconstructive mechanism applied to a recoverable event ordering. The next chapter derives that ordering itself.

Chapter 27

Memory Before Time

This is the most radical inversion. Standard accounts place time first and locate memory within time. The MEM|8 position is that temporal structure is itself derived from recoverable event orderings.

Definition 27.1 (Recoverable Event Ordering)**.** A *recoverable event ordering* over a set of events $\mathcal{E}$ is a partial order $\prec$ on $\mathcal{E}$ such that, for any admissible reconstruction operator R and any $e_i \prec e_j$, the distinction associated with e_i can be recovered from the historical state $H_{e_j} = \{e_k : e_k \prec e_j\}$.

Theorem 27.2 (Time Emerges from Recoverable Order)**.** *Temporal structure, in the operational sense relevant to memory, learning, and cognition, is equivalent to recoverable event ordering. Without reconstructable event relations, temporal distinctions are operationally undefined.*

Proof. The direction *temporal order implies recoverable order*: if e_i causally precedes e_j in Minkowski spacetime, then H_{e_j} contains information traceable to e_i, and an admissible R can in principle recover the relevant distinction.

The direction *recoverable order implies temporal structure*: if for all $e_i \prec e_j$ the distinction of e_i is recoverable from H_{e_j}, then $\prec$ constitutes a causal precedence relation. The set $(\mathcal{E}, \prec)$ with this property is a *causal set* in the sense of Bombelli et al.; causal sets are known to encode the topological and differential structure of spacetime in the continuum limit. Therefore recoverable event ordering is sufficient to reconstruct operational temporal structure. □

Corollary 27.3. *Memory does not occur* in *time as a background stage. Time is the structure that emerges from the recoverable ordering of events in a memory field. The arrow of time is the direction in which reconstruction of earlier events from later states is possible.*

We arrive at the foundational statements of MEM|8:

$$\boxed{\text{A memory is not a preserved past.}}$$

$$\boxed{\text{A memory is the continued ability to regenerate distinctions from transformation.}}$$

$$\boxed{\text{Memory} = \text{recoverable resonant structure maintained by flux.}}$$

Bibliography: Neuroscience, Biology, and Cognitive Science

Bibliography

[1] Anderson, J. R. (1983). *The Architecture of Cognition*. Harvard University Press.

[2] Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review*, 111(4), 1036–1060.

[3] Amit, D. J. (1989). *Modeling Brain Function: The World of Attractor Neural Networks*. Cambridge University Press.

[4] Buzsáki, G. (2006). *Rhythms of the Brain*. Oxford University Press.

[5] Cabral, J., Fernandes, F. F., & Shemesh, N. (2023). Intrinsic macroscale oscillatory modes driving long range functional connectivity in female rat brains detected by ultrafast fMRI. *Nature Communications*, 14, 375.

[6] Chen, J., Ige, A. S., Runge, K., Deymier, P. A., & Yan, X. (2026). Topological acoustic synapse for high-dimensional neuromorphic computing. *Science Advances*, 12(24), eaec6633.

[7] Cohen, L. D., et al. (2013). Metabolic turnover of synaptic proteins. *Science*, 342(6161), 1094–1097.

[8] Dayan, P., & Abbott, L. F. (2001). *Theoretical Neuroscience*. MIT Press.

[9] Deco, G., Jirsa, V. K., & McIntosh, A. R. (2011). Emerging concepts for brain dynamics. *Nature Reviews Neuroscience*, 12(1), 43–56.

[10] Fornasiero, E. F., et al. (2018). Precisely measured protein lifetimes. *Nature Communications*, 9, 4741.

[11] Friston, K. (2010). The free-energy principle. *Nature Reviews Neuroscience*, 11(2), 127–138.

[12] Gramlich, M. W., & Klyachko, V. A. (2021). Motor-mediated regulation of synaptic vesicle mobility. *Biophysical Journal*, 120(4), 567–582.

[13] Hebb, D. O. (1949). *The Organization of Behavior*. Wiley.

[14] Hirokawa, N., Noda, Y., Tanaka, Y., & Niwa, S. (2009). Kinesin superfamily motors. *Nature Reviews Molecular Cell Biology*, 10(10), 682–696.

[15] Hopfield, J. J. (1982). Neural networks and physical systems. *Proceedings of the National Academy of Sciences*, 79(8), 2554–2558.

[16] Joensuu, M., et al. (2016). Subdiffractional tracking of synaptic vesicles. *Nature Communications*, 7, 11714.

[17] Kelso, J. A. S. (1995). *Dynamic Patterns*. MIT Press.

[18] Parkes, M., Landers, N. L., & Gramlich, M. W. (2023). Recently recycled synaptic vesicles use multi-cytoskeletal transport and differential presynaptic capture probability to establish a retrograde net flux during ISVE in central neurons. *Frontiers in Cell and Developmental Biology*, 11, 1286915.

[19] Rizzoli, S. O. (2014). Synaptic vesicle recycling. *Nature Reviews Neuroscience*, 15(10), 658–673.

[20] Rolls, E. T. (2016). *Cerebral Cortex: Principles of Operation*. Oxford University Press.

[21] Sporns, O. (2011). *Networks of the Brain*. MIT Press.

[22] Truckenbrodt, S., et al. (2018). Newly produced synaptic vesicle proteins. *Science*, 359(6371), 157–161.

[23] Tononi, G. (2008). Consciousness as integrated information. *Biological Bulletin*, 215(3), 216–242.

[24] Chenoweth, C., & Chenoweth, A. (2025). *MEM|8: A wave-based cognitive architecture for multimodal memory integration and consciousness simulation*. Zenodo. `https://doi.org/10.5281/zenodo.16436297`

Bibliography: Mathematics, Physics, and Dynamical Systems

Bibliography

[1] Arnold, V. I. (1989). *Mathematical Methods of Classical Mechanics* (2nd ed.). Springer.

[2] Arnold, V. I. (2006). *Ordinary Differential Equations*. Springer.

[3] Ash, R. B. (1965). *Information Theory*. Dover.

[4] Billingsley, P. (1995). *Probability and Measure* (3rd ed.). Wiley.

[5] Bowen, R. (1978). *Entropy for Group Endomorphisms and Homogeneous Spaces*. Springer.

[6] Cover, T. M., & Thomas, J. A. (2006). *Elements of Information Theory* (2nd ed.). Wiley.

[7] Crutchfield, J. P. (1994). The calculi of emergence. *Physica D*, 75(1–3), 11–54.

[8] Evans, L. C. (2010). *Partial Differential Equations* (2nd ed.). American Mathematical Society.

[9] Ghrist, R. (2014). *Elementary Applied Topology*. Createspace.

[10] Guckenheimer, J., & Holmes, P. (1983). *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*. Springer.

[11] Hatcher, A. (2002). *Algebraic Topology*. Cambridge University Press.

[12] Horn, R. A., & Johnson, C. R. (2013). *Matrix Analysis* (2nd ed.). Cambridge University Press.

[13] Jaynes, E. T. (2003). *Probability Theory: The Logic of Science*. Cambridge University Press.

[14] Katok, A., & Hasselblatt, B. (1995). *Introduction to the Modern Theory of Dynamical Systems*. Cambridge University Press.

[15] Kreyszig, E. (1989). *Introductory Functional Analysis with Applications*. Wiley.

[16] Mac Lane, S. (1998). *Categories for the Working Mathematician* (2nd ed.). Springer.

[17] Mézard, M., & Montanari, A. (2009). *Information, Physics, and Computation*. Oxford University Press.

[18] Nakahara, M. (2003). *Geometry, Topology and Physics* (2nd ed.). Taylor & Francis.

[19] Pikovsky, A., Rosenblum, M., & Kurths, J. (2001). *Synchronization: A Universal Concept in Nonlinear Sciences*. Cambridge University Press.

[20] Reed, M., & Simon, B. (1972). *Methods of Modern Mathematical Physics I: Functional Analysis*. Academic Press.

[21] Reed, M., & Simon, B. (1975). *Methods of Modern Mathematical Physics II: Fourier Analysis, Self-Adjointness*. Academic Press.

[22] Rudin, W. (1991). *Functional Analysis* (2nd ed.). McGraw-Hill.

[23] Ruelle, D. (1989). *Chaotic Evolution and Strange Attractors*. Cambridge University Press.

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

[25] Smale, S. (1967). Differentiable dynamical systems. *Bulletin of the American Mathematical Society*, 73(6), 747–817.

[26] Strogatz, S. H. (2018). *Nonlinear Dynamics and Chaos* (2nd ed.). Westview Press.

[27] Takens, F. (1981). Detecting strange attractors in turbulence. In D. Rand & L.-S. Young (Eds.), *Dynamical Systems and Turbulence*. Springer.

[28] Temam, R. (1997). *Infinite-Dimensional Dynamical Systems in Mechanics and Physics* (2nd ed.). Springer.

[29] Tu, L. W. (2011). *An Introduction to Manifolds* (2nd ed.). Springer.

[30] von Neumann, J. (1955). *Mathematical Foundations of Quantum Mechanics*. Princeton University Press.

[31] Wiggins, S. (2003). *Introduction to Applied Nonlinear Dynamical Systems and Chaos* (2nd ed.). Springer.

[32] Zeidler, E. (1990). *Applied Functional Analysis*. Springer.