

Synthetic Reification and Epistemic Constraint Deformation

Coherence-Driven Belief Formation and Its Vulnerability to Unverifiable Synthetic Content

Flyxion

Independent Researcher

May 4, 2026

Abstract

Cognitive reconstruction prioritizes coherence over provenance: the mind accepts representations not because they can be verified against lived experience but because they satisfy internal consistency constraints. This paper formalizes that vulnerability through the CLIO architecture (Cheng, Broadbent, and Chappell, 2025) and the RSVP field framework, demonstrating that any reconstructive system with non-invertible projection and coherence-based acceptance is structurally indifferent to the origin of its inputs. The vulnerability has particular force for content that cannot be verified through direct experience: news events, scientific findings, historical sequences, social norms, and demonstrations of tool use or physical phenomena. When such content is rendered in high-fidelity synthetic form — visually and aurally concrete, narratively consistent, and perceptually dense — it satisfies the same acceptance conditions as genuine testimony or observation. We develop a stability theorem governing when repeated exposure to synthetic reifications of unverifiable content produces stable shifts in the constraint surface governing belief and trust, and when those shifts become self-reinforcing. The implications are especially significant for developmental contexts, where trust calibration and epistemic norms are being established for the first time against a corpus that now includes arbitrary synthetic content.

1. Introduction

Human cognition depends extensively on representations of things that cannot be directly experienced. We learn physics through diagrams, demonstrations, and simulations. We form beliefs about historical events through testimony, documentary footage, and narrative reconstruction. We acquire social norms, trust institutions, and calibrate our sense of what is normal and possible largely through mediated content rather than first-hand encounter. This is not a deficiency but a structural feature of cognition at scale: no individual can verify more than a small fraction of what they know through direct experience, so the cognitive system must rely on representations whose provenance it cannot fully assess.

The vulnerability this creates has always existed. The availability heuristic — the tendency to weight the probability and importance of events by how easily they can be imagined or recalled — means that vivid, concrete representations carry disproportionate epistemic weight relative to their evidential value. Imagery and sound make abstract claims tractable, which is why mnemonic systems, scientific visualization, and pedagogical demonstration are powerful tools. But the same mechanism that makes a physics simulation educationally effective also makes a fabricated news event cognitively compelling. When arbitrary actors can produce high-fidelity synthetic reifications of things no one can directly verify — alternate historical sequences, contested scientific findings, demonstrations of social norms, fabricated testimony — the epistemic cost of that availability becomes substantial.

Standard accounts treat this as a problem of deception detection: can people tell real from synthetic? The present work advances a stronger claim. Any reconstructive system satisfying three structural properties — non-invertible projection from state space to representation space, selection over admissible trajectories rather than retrieval of a fixed trace, and acceptance governed by coherence rather than provenance — is necessarily equivalent to a CLIO-type descent operator. The question of whether a representation is real is not the relevant question for such a system. The relevant question is whether it is admissible: whether it satisfies the coherence constraints the system applies. Synthetic reifications of unverifiable content that are perceptually dense, narratively consistent, and structurally proximal to existing belief structures will satisfy those constraints regardless of their origin. The central claim of this paper is a necessity result: given those three properties, systematic deformation of the constraint surface

governing belief and trust under sufficiently coherent, repeated synthetic inputs is not merely possible but unavoidable.

The concern is amplified in developmental contexts. Children and individuals encountering a domain for the first time have no prior verified experience against which to compare synthetic representations. The trust calibration and epistemic norms that develop through learning are being established against a corpus that now includes arbitrary synthetic content, and the cognitive machinery for distinguishing reliable from unreliable sources is itself shaped by the inputs it receives during formation.

We formalize these claims within the RSVP field framework. The argument proceeds in stages: Section 2 establishes the CLIO equivalence class and its architectural vulnerability; Section 3 develops the RSVP formalization including trajectory admissibility and projection non-invertibility; Section 4 derives the constraint surface stability theorem, which is the central result; Section 5 identifies the specific features of synthetic reification that satisfy the theorem's deformation conditions; Section 6 addresses the role of social correction as external constraint enforcement and its failure modes; and Section 7 discusses the limits of the claim and open empirical questions.

2. Memory Reconsolidation as Recursive Confidence Resolution

2.1. The Reconstructive Baseline

The contemporary scientific consensus holds that autobiographical memory is reconstructive rather than reproductive. Foundational work by Elizabeth Loftus, Marcia K. Johnson, and Daniel Schacter demonstrates that recall involves the reactivation and recombination of distributed neural representations rather than the retrieval of a stable stored record.

Within this framework, memory errors arise not from corruption of a stored trace but from the dynamics of reconstruction itself. The central vulnerability is the source monitoring problem: the system must determine whether a reactivated pattern originated from perception, imagination, or inference, yet no intrinsic tag encodes this provenance. Instead, the system relies on heuristic cues — perceptual richness, contextual embedding, and retrieval fluency — to infer origin.

These cues are neither necessary nor sufficient indicators of truth. Empirical

work shows that imagined events can acquire perceptual detail through rehearsal, while genuine memories may degrade or fragment over time. The result is a regime in which true and false memories are generated by the same underlying processes and differ primarily in their causal history, not in their structural form at recall.

A clarification is necessary at this point regarding what is meant by perceptual richness in the reconstruction process. The framework developed here does not presuppose that cognitive reconstruction involves pictorial mental imagery in the sense of visualizing scenes as if viewing them on an internal screen. The phenomenology of mental imagery varies substantially across individuals: research on aphantasia and hyperphantasia suggests that roughly two to three percent of the population experience highly vivid, picture-like mental imagery, a similar proportion experience little or none, and the majority occupy an intermediate range where imagery is present in varying degrees of vividness and modality [1]. The cues that govern reconstruction — perceptual richness, contextual embedding, retrieval fluency — are therefore not reducible to visual vividness. They are properties of the constraint structure: the density and coherence of the local sections $\mathcal{F}_n(U_i)$ that a given input populates. A highly coherent auditory narrative, a kinesthetic demonstration, or a spatially consistent description can each populate constraint patches and reduce the relevant cocycles without requiring any pictorial phenomenology on the part of the subject. The vulnerability this paper describes is not a vulnerability of those who visualize vividly; it is a structural property of any system that reconstructs under compression and resolves under coherence, regardless of the phenomenal character of that reconstruction.

Standard accounts treat this as a limitation of memory: reconstruction introduces distortion as a side effect of an otherwise efficient system. The present work advances a stronger claim. Reconstruction is not merely a process that can produce errors; it implements a specific class of iterative resolution algorithm whose acceptance criterion is coherence rather than provenance.

Three bodies of work ground this claim empirically and sharpen the mechanism. Tulving’s ephory principle holds that retrieval is not direct access to a stored trace but the product of an interaction between a retrieval cue and the engram: what is recovered is always a joint construction, and its content is shaped as much by the cue as by the original encoding [2]. This makes the retrieval cue a structural input to the reconstruction process, not merely a pointer to a stored location. The admissible neighborhood $\mathcal{N}(s_n)$ in the CLIO formalization is pre-

cisely the set of states reachable from the cue under the current constraint configuration: ecphory defines which states are locally accessible. Fauconnier and Turner’s theory of conceptual blending identifies the mechanism by which cues recruit material from multiple mental spaces and integrate them into emergent structures that were not present in any source [3]. Blending is not error but the standard operation of conceptual combination, and its output is a coherent but novel configuration rather than a reproduction of prior content. In the present framework, blending corresponds to the branching exploration phase in which candidate states are assembled from partial matches across constraint patches; the resulting state satisfies local coherence conditions without necessarily corresponding to any single prior trajectory. Semantic priming research demonstrates that the activation of one concept systematically lowers the threshold for activating related concepts [4], which in the reconstruction context means that a retrieval cue does not merely identify a target but reshapes the landscape of accessible states. Priming is the empirical signature of the neighborhood structure $\mathcal{N}(s_n)$: it determines which candidate reconstructions are energetically available at the moment of retrieval. Together, ecphory, blending, and priming establish that the reconstruction process is generative, cue-sensitive, and structurally indifferent to provenance at every stage, from initial cue activation through candidate assembly to final acceptance.

2.2. CLIO as a Computational Model of Reconsolidation

We propose that memory reconsolidation can be modeled as a recursive confidence-resolution process structurally equivalent to the CLIO architecture [5]. CLIO operates by iteratively refining a semantic state through branching exploration and constraint satisfaction until a confidence threshold is reached. The system evolves a state s through successive updates,

$$s_{n+1} = \mathcal{R}(s_n, \mathcal{N}(s_n)),$$

where $\mathcal{N}(s_n)$ denotes the neighborhood of candidate continuations and \mathcal{R} is a resolution operator selecting among them. Each candidate state is assigned a confidence value $c(s)$, and the process terminates when $c(s) \geq \tau$ for some threshold τ , at which point the state is accepted as resolved.

The mapping onto reconsolidation is as follows. Retrieval cue triggers partial reactivation of a memory trace, corresponding to initialization of the semantic state s_0 . Associative processes then recruit related perceptual, semantic, and

affective elements, generating candidate reconstructions that correspond to the branching exploration phase. These candidates are evaluated for internal coherence and compatibility with existing constraints, corresponding to the resolution step. Reconstruction stabilizes when a sufficiently coherent configuration is achieved, satisfying the termination condition $c(s) \geq \tau$.

Under this mapping, reconsolidation is not the retrieval of a fixed trace but the convergence of a search process over a space of admissible reconstructions. The crucial structural property — and the point of vulnerability — is that acceptance depends on confidence rather than provenance. The system does not verify whether a candidate state corresponds to an actual past trajectory; it evaluates only whether the state satisfies coherence constraints to a sufficient degree.

Empirical findings in reconsolidation research support this interpretation. Successful reconsolidation is associated with reduction in prediction error and stabilization of neural activation patterns, while failed reconsolidation exhibits oscillatory or unstable dynamics [9, 10]. This aligns directly with CLIO’s characteristic convergence behavior: correct resolutions exhibit a negative uncertainty gradient converging on a stable state, while incorrect or unresolvable states show oscillation and positive gradient without convergence [5]. The computational signature of successful reconsolidation and successful CLIO resolution fall within the same equivalence class: both exhibit convergence under a negative uncertainty gradient and instability under unresolved constraint conflict [5].

2.3. Architectural Vulnerability: Coherence Without Provenance

Standard accounts of false memory emphasize post-hoc source confusion: the system reconstructs an event and subsequently misattributes its origin. The CLIO-based account shifts the locus of explanation. False memories arise not because provenance is misidentified after reconstruction, but because provenance is never part of the acceptance criterion.

Formally, let s denote a reconstructed memory state. The system evaluates

$$c(s) = f(\text{coherence, consistency, fluency, } \dots)$$

and accepts s when $c(s) \geq \tau$. There is no term in $c(s)$ corresponding to verification of origin in lived experience. As a result, any input that produces a candidate state satisfying the coherence constraints will be accepted regardless of its causal source.

This implies that synthetic perceptual inputs — such as those generated by contemporary generative video systems — are not merely misleading stimuli but structurally valid inputs to the reconstruction algorithm. If such inputs produce candidate states with sufficient perceptual richness, spatial coherence, and narrative compatibility, they will pass the same acceptance threshold as genuine memories. The consequence is a structural vulnerability rather than an edge case: the system is optimized to resolve ambiguity through coherence, but coherence is not a reliable indicator of truth when high-fidelity synthetic inputs are available.

This prediction goes beyond standard reconstructive accounts. It implies that increasing the perceptual fidelity and repeatability of synthetic inputs should systematically increase their probability of acceptance, even when their artificial origin is known to the subject. False memories under this account are not anomalous failures but expected outputs of an architecture that resolves for admissibility rather than provenance. To formalize this acceptance mechanism and its vulnerability, we require a representation in which reconstructed states, their admissibility conditions, and the absence of provenance constraints can be expressed explicitly; this is provided by the RSVP field framework introduced in the next section.

2.4. Minimal Operator Form of Reconsolidation Dynamics

The CLIO mapping admits a minimal operator characterization. Let \mathcal{S} denote the space of candidate reconstruction states and $c : \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$ a coherence functional. Define the resolution operator

$$\mathcal{R}(s) = \arg \max_{s' \in \mathcal{N}(s)} c(s'),$$

where $\mathcal{N}(s)$ is the admissible neighborhood of continuations from s . Iterated application produces the discrete dynamical system $s_{n+1} = \mathcal{R}(s_n)$, with fixed points satisfying $s^* = \mathcal{R}(s^*)$ corresponding to locally maximal coherence states. Setting $U(s) = -c(s)$, the dynamics implement a descent process $U(s_{n+1}) \leq U(s_n)$; successful reconsolidation is convergence to s^* , and failure is non-convergent or oscillatory behavior.

Crucially, $c(s)$ depends only on internal structural properties of s . Let $\mathcal{S}_{\text{real}} \subset \mathcal{S}$ denote states corresponding to lived trajectories and $\mathcal{S}_{\text{synthetic}}$ denote those gen-

erated without field-space provenance. Then

$$\mathcal{R} : \mathcal{S} \rightarrow \mathcal{S} \quad \text{with} \quad \mathcal{R}(\mathcal{S}_{\text{real}} \cup \mathcal{S}_{\text{synthetic}}) \subset \mathcal{S}.$$

The operator \mathcal{R} is indifferent to provenance: false memories are not misclassifications but fixed points of the same descent process defined over an enlarged admissible space. This operator-level indifference to provenance is the mechanism that, under non-invertible projection, lifts from state dynamics to constraint surface deformation.

Induced dynamics on constraint space. Let $\pi : X \rightarrow M$ be the projection defined in Section 3 and let \mathcal{C} denote the space of admissibility sheaves over X . The CLIO operator $\mathcal{R} : \mathcal{S} \rightarrow \mathcal{S}$ induces an operator on constraint configurations via pushforward along π :

$$\mathcal{T}(C) = \text{CLIO}(C) = \pi_* \circ \mathcal{R} \circ \pi^{-1}(C),$$

defined on equivalence classes of trajectories under π . Because π is non-invertible, $\pi^{-1}(C)$ is set-valued, and \mathcal{T} is correspondingly a set-valued descent operator on \mathcal{C} . The constraint dynamics of Section 4 therefore arise as the pushforward of the state-level CLIO operator under non-invertible projection. In particular, the indifference of \mathcal{R} to provenance lifts to \mathcal{T} as indifference to the origin of admissible trajectories, implying that deformation of the constraint surface is a direct consequence of the state-level acceptance criterion rather than an independent assumption. The entire paper describes a single dynamical system viewed at three levels of resolution: state reconstruction (\mathcal{R} on \mathcal{S}), constraint evolution (\mathcal{T} on \mathcal{C}), and long-time flow (\bar{V} on the hybrid system).

3. The RSVP Formalization

3.1. Memory as Trajectory in a Field

We model autobiographical memory not as a discrete stored object but as a trajectory within a continuous state space. Let the state of the system be given by a field

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

where Φ is a scalar potential encoding salience and meaning, \mathbf{v} is a vector field encoding directional structure and temporal sequencing, and S is an entropy

field encoding uncertainty over admissible continuations.

A lived experience corresponds to a trajectory γ traversed through this field. Crucially, recall does not retrieve γ as a stored object; instead, it reconstructs a trajectory $\hat{\gamma}$ that satisfies local and global constraints induced by the current field configuration. This establishes a first structural property: reconstruction is generative rather than reproductive. The system solves for a trajectory compatible with constraints rather than retrieving a uniquely determined past path.

3.2. Projection and Non-Invertibility

Let π denote a projection from field space X to model space M , where M consists of consciously accessible, narratively structured states. The projection π is necessarily many-to-one and lossy: multiple distinct trajectories in X can produce indistinguishable representations in M .

It follows that π is non-invertible. There exists no mapping π^{-1} such that, given a model-space element $m \in M$, the originating trajectory in X can be uniquely recovered. Formally, for any $m \in M$, the preimage

$$\pi^{-1}(m) = \{ \gamma \subset X \mid \pi(\gamma) = m \}$$

is a set of equivalence-class trajectories rather than a singleton.

This non-invertibility is not a limitation of implementation but a structural feature of any compressed representational system. It implies that the system cannot, even in principle, determine whether a given model-space state originated from a particular trajectory in field space. The absence of a provenance check is therefore not an oversight but a mathematical consequence of the projection structure.

3.3. Admissibility Without Provenance

Given the absence of invertibility, reconstruction must operate by selecting among candidate trajectories in X that are compatible with a given model-space state. Let \mathcal{A} denote the set of admissible trajectories under current constraints. A reconstructed trajectory $\hat{\gamma}$ is accepted when

$$\hat{\gamma} \in \mathcal{A},$$

where admissibility is defined by coherence, consistency with existing structure, and compatibility with constraint conditions. Provenance does not appear in the definition of \mathcal{A} .

We distinguish two classes of trajectories. A real trajectory γ satisfies $\gamma \in \mathcal{A}_{\text{real}}$, meaning it was physically traversed. A synthetic trajectory $\tilde{\gamma}$ may satisfy $\tilde{\gamma} \in \mathcal{A}$ while $\tilde{\gamma} \notin \mathcal{A}_{\text{real}}$. Because admissibility is evaluated without reference to provenance, both classes are treated identically by the reconstruction process. Under projection, if

$$\pi(\tilde{\gamma}) \approx \pi(\gamma),$$

then the system cannot distinguish between them. The acceptance criterion is equivalence under projection, not correspondence in field space.

3.4. Iterated Reconsolidation and Entropy Reduction

We model reconsolidation as an iterative process

$$C_{n+1} = \text{CLIO}(C_n),$$

where C_n denotes the current constraint configuration. Introduction of a synthetic trajectory $\tilde{\gamma}$ yields

$$C_{n+1} = \text{CLIO}(C_n \cup \tilde{\gamma}).$$

Each iteration reduces entropy $S(\tilde{\gamma})$ over the accepted trajectory by increasing its coherence with surrounding structure. Repeated exposure produces

$$S(\tilde{\gamma}) \rightarrow 0,$$

driving the trajectory toward maximal admissibility. This captures the illusory truth effect as a dynamical process: repetition reduces uncertainty, increasing the probability of acceptance independently of provenance.

The direction of the entropy gradient is diagnostically significant. Where $\tilde{\gamma}$ is locally admissible, the CLIO resolution loop converges with negative entropy gradient — the same signature as successful reconsolidation of a real memory. Where $\tilde{\gamma}$ violates constraints, the loop exhibits oscillation and positive gradient, corresponding to failed reconsolidation and eventual rejection. The admissibility of the synthetic input therefore determines not only whether it is accepted but how it is accepted: indistinguishably from a genuine memory trace.

3.5. Structural Consequence

The combination of non-invertible projection and admissibility-based selection yields a decisive consequence. Model-space states can be generated that are indistinguishable under π from those produced by real trajectories, and once admitted into the reconstruction process, these states are stabilized through iterative entropy reduction. The system therefore admits elements of M that correspond to no trajectory in $\mathcal{A}_{\text{real}}$ but are treated as though they did.

This is not a peripheral limitation of memory but a defining feature of the architecture. Any system that reconstructs under compression and resolves under coherence will exhibit this vulnerability. The question is not whether such a system can be deceived by high-fidelity synthetic inputs but under what conditions the deception stabilizes and what its cumulative effect on the constraint surface constituting identity will be. We address this in the following section.

4. The Constraint Surface and Its Deformation

This section constructs the formal apparatus required for the stability theorem. It proceeds in three stages: first, the sheaf formulation of the constraint surface and identity; second, the obstruction functional measuring failure of admissible gluing; third, the stability theorem itself, which is the central result of the paper. A temporal extension follows, upgrading the discrete-time theorem to a hybrid dynamical system with an explicit phase-transition condition.

Proposition (Lyapunov-entropy identification). *Let $V(C) = \|\text{Obs}\|^2 + S$ be the coherence functional on the space of admissibility sheaves. The entropy term $S(\tilde{\gamma})$ associated with a candidate trajectory $\tilde{\gamma}$ is the variational derivative of V along trajectory directions:*

$$S(\tilde{\gamma}) = \frac{\delta V}{\delta \tilde{\gamma}}.$$

Consequently, V generates both the constraint surface evolution and the entropy dynamics on trajectories within a single potential. The CLIO descent $V(\mathcal{T}(C)) \leq V(C)$ and the entropy reduction $S_{n+1}(\tilde{\gamma}) \leq S_n(\tilde{\gamma})$ under repeated exposure are both manifestations of descent under the same functional, related by the identification above.

This identification means the entire system is a single Lyapunov-driven field theory on constraint space: one potential, two views of its descent. The hybrid dynamics and the phase transition condition introduced later in this section are direct consequences.

4.1. Identity as a Constraint Surface

We define the autobiographical system at time n as a constraint configuration C_n , encoding the accumulated structure of admissible reconstructions. The constraint surface Σ_n is the manifold of trajectories admissible under C_n :

$$\Sigma_n = \{ \gamma \subset X \mid \gamma \in \mathcal{A}(C_n) \}.$$

Equivalently, $\Sigma_n = \Gamma(X, \mathcal{F}_{C_n})$ is the space of global sections of the admissibility sheaf: identity is not a collection of stored episodes but the global section space of a sheaf whose local data encodes what can be reconstructed as one's own. Deformation of identity corresponds to a morphism of sheaves $\mathcal{F}_{C_n} \rightarrow \mathcal{F}_{C_{n+1}}$ induced by admitted synthetic trajectories, which continuously modifies the sheaf rather than replacing it.

Σ_n is not a collection of stored memories. It is the structure that governs which trajectories can be reconstructed as one's own: which experiences feel continuous with the past, which self-attributions are available, which future projections are coherent with prior identity. Identity, under this formalization, is not constituted by the trajectories themselves but by the constraint surface that bounds and shapes their reconstruction.

CLIO convergence corresponds to convergence of Σ_n . When iterative reconsolidation stabilizes — that is, when successive applications of the resolution operator produce no further change in the constraint configuration — a stable surface Σ^* is induced. This is the equilibrium condition for autobiographical identity: not a fixed set of memories but a stable set of admissibility conditions under which memory can be reconstructed.

4.2. Synthetic Injection as Perturbation

Let $\tilde{\gamma} \in \mathcal{A}(C_n)$ but $\tilde{\gamma} \notin \mathcal{A}_{\text{real}}$: a synthetic trajectory that satisfies current admissibility conditions without corresponding to a lived experience. Its introduction into the constraint system yields the update

$$C_{n+1} = \text{CLIO}(C_n \cup \tilde{\gamma}).$$

We define this as a perturbation operator $\mathcal{P}_{\tilde{\gamma}}$ acting on the constraint configuration:

$$\mathcal{P}_{\tilde{\gamma}}(C_n) := \text{CLIO}(C_n \cup \tilde{\gamma}).$$

Two cases must be distinguished. If $\tilde{\gamma} \notin \mathcal{A}(C_n)$, the CLIO resolution loop exhibits high entropy and oscillatory behavior; the trajectory is not integrated and $C_{n+1} \approx C_n$. If $\tilde{\gamma} \in \mathcal{A}(C_n)$, the loop converges with negative entropy gradient; the trajectory enters the constraint update and C_{n+1} differs from C_n by the contribution of $\tilde{\gamma}$. Deformation of Σ occurs only through admitted perturbations. The admissibility condition is therefore the gate through which synthetic inputs must pass before any structural effect is possible.

4.3. Local Admissibility and Obstruction

Not all admissible trajectories integrate with equal ease. We distinguish local from global admissibility using the language of sheaf cohomology, which provides the precise geometric backbone for the stability theorem that follows. This formulation is minimal in the sense that it encodes precisely the local-to-global gluing problem inherent in autobiographical reconstruction, without introducing additional structure beyond what is required to formalize admissibility.

Let $\mathfrak{U} = \{U_i\}$ be an open cover of X , where each U_i represents a local constraint patch: a temporal window, thematic cluster, sensory modality, or coherent episodic neighborhood. The constraint configuration C_n defines a sheaf of admissible sections \mathcal{F}_n on X : for each open U_i , the set $\mathcal{F}_n(U_i)$ consists of trajectory pieces locally compatible with the constraints in C_n , with restriction maps encoding coherent overlap. The sheaf property encodes the gluing axioms for autobiographical reconstruction.

A synthetic trajectory $\tilde{\gamma}$ provides local data: for each U_i , a section $s_i \in \mathcal{F}_n(U_i)$. Local admissibility is the condition that such sections exist patch by patch. Global admissibility is the stronger condition that these local sections glue to a global section $s \in \Gamma(X, \mathcal{F}_n)$. The obstruction to gluing is measured by the first Čech cohomology group $\check{H}^1(\mathfrak{U}, \mathcal{F}_n)$.

Given local sections $\{s_i\}$, define the Čech 1-cocycle measuring failure of gluing on pairwise overlaps:

$$\omega_{ij} = s_i|_{U_i \cap U_j} - s_j|_{U_i \cap U_j} \in \mathcal{F}_n(U_i \cap U_j).$$

The class $[\omega] \in \check{H}^1(\mathfrak{U}, \mathcal{F}_n)$ vanishes if and only if the local sections can be adjusted to glue consistently into a global section. We define the scalar obstruction functional as

$$\text{Obs}(\tilde{\gamma}, C_n) := \inf_{\text{refinements}} \left\| [\omega] \right\|_{\check{H}^1(\mathfrak{U}, \mathcal{F}_n)},$$

where the norm is taken in a suitable topology on the cochain complex and the infimum is taken over all refinements of \mathfrak{U} , converging in the direct limit to the Čech cohomology of X . This infimum is well-defined when the cochain complex carries a Banach or Hilbert structure compatible with the trajectory space metric, and equals zero if and only if $[\omega]$ is trivial in the limit sheaf cohomology. When $\text{Obs}(\tilde{\gamma}, C_n) = 0$, the synthetic trajectory is globally admissible and integrates without constraint restructuring. When $\text{Obs}(\tilde{\gamma}, C_n) \approx 0$, it is glueable with small deformation of \mathcal{F}_n . When $\text{Obs}(\tilde{\gamma}, C_n) \gg 0$, a nontrivial cohomology class prevents integration without large-scale reorganization of the constraint surface.

This formalization makes the distinction between local and global admissibility precise. Local admissibility corresponds to the existence of the 0-cochains $\{s_i\}$. Global admissibility requires the cohomology class $[\omega]$ to be trivial. A trajectory can be locally plausible — passing through each constraint patch individually — while generating a nontrivial obstruction class that prevents global integration. The danger zone is the interior of admissibility: inputs whose cocycles are small and can be killed at low entropy cost. These are the synthetic trajectories that pass local constraint checks, generate near-zero obstruction, and integrate without triggering the reorganization that would make their presence detectable.

4.4. Stability Theorem for the Constraint Surface

Theorem (Constraint Surface Stability Under Synthetic Injection). *Let $\Sigma_n = \Gamma(X, \mathcal{F}_n)$ be the space of global sections of the admissibility sheaf. Consider a sequence of synthetic perturbations $\{\tilde{\gamma}_k\}$. Then:*

(i) *If $\text{Obs}(\tilde{\gamma}_k, C_n) \geq \delta > 0$ for all k , then $[\omega_k] \neq 0$ in \check{H}^1 and the surface is stable: $\Sigma_{n+k} \approx \Sigma_n$.*

(ii) *If $\text{Obs}(\tilde{\gamma}_k, C_n) \leq \delta$ but the minimal entropy cost to kill $[\omega_k]$ satisfies $S(\tilde{\gamma}_k | C_n) \geq \varepsilon > 0$, perturbations are admitted but deformation remains bounded:*

$$d(\Sigma_{n+k}, \Sigma_n) \leq M\varepsilon$$

for some constant M depending on the constraint structure.

(iii) *If there exists a subsequence such that $\text{Obs}(\tilde{\gamma}_{k_j}, C_n) \rightarrow 0$ and $S(\tilde{\gamma}_{k_j} | C_n) \rightarrow 0$, then the sheaf \mathcal{F}_n deforms continuously in its moduli space and*

$$\Sigma_n \rightarrow \Sigma'_n \neq \Sigma_n.$$

Proof sketch. Define a coherence functional $V(C)$ as the sum of squared obstruction norms plus entropy terms across the constraint surface: $V(C) = \|\text{Obs}\|^2 + S$, where $S(\tilde{\gamma})$ is understood as the variational derivative of V along trajectory directions, $S(\tilde{\gamma}) = \delta V / \delta \tilde{\gamma}$, so that V generates both the constraint evolution and the entropy dynamics on trajectories within a single potential. The existence of such a functional is guaranteed by the convergence properties established in Section 2: CLIO dynamics define a descent process over a coherence landscape, and V is the induced potential on that landscape. CLIO acts as a descent procedure reducing V toward a local minimum corresponding to a stable sheaf configuration \mathcal{F}^* and induced surface Σ^* . For case (i), nontrivial cohomology classes cannot be killed without finite cost; V does not decrease at the putative new minimum and perturbations are rejected. For case (ii), the entropy lower bound ε limits how far the minimum of V can shift per step; the surface deforms but the displacement is controlled by $M\varepsilon$. In this regime, deformation remains within the basin of attraction of the original equilibrium, and removal of the perturbation sequence allows recovery under further CLIO iteration. For case (iii), vanishing obstruction and vanishing entropy cost mean each perturbation contributes negligibly to V pointwise, but the sequence is not summable in the minimizer space even though it is pointwise small in V : this is a non-uniform convergence phenomenon in which the minimizer drifts even as individual increments vanish. Perturbations are not merely added to the configuration but used to re-minimize the functional, so successive equilibria shift incrementally and the cumulative displacement is finite and nonzero. The sheaf \mathcal{F}_n moves continuously through its moduli space, the induced surface Σ_n drifts without bound from its initial configuration, and because each new configuration is itself a CLIO-stable minimum of V , the system cannot recover the original surface through iteration alone.

4.5. Temporal Relaxation of the Coherence Functional

The stability theorem as stated treats reconsolidation as a purely discrete iterative process, indexed by retrieval events without reference to the dynamics of the system between events. This is a structural limitation: in biological systems, constraints relax, decay, and reorganize in the absence of retrieval, and the entropy associated with a given trace may increase between exposures as the system drifts away from a previously stabilized configuration. To address this, we extend the model to a hybrid dynamical system in which continuous relaxation is punctuated by discrete CLIO resolution events.

Between retrieval events, the constraint configuration evolves under a continuous-time dissipative flow:

$$\frac{dC}{dt} = -\lambda \nabla V(C) + \eta(t),$$

where $\lambda > 0$ is a relaxation constant governing the rate at which the system returns toward prior constraint configurations, and $\eta(t)$ captures stochastic reorganization, sleep-driven consolidation, and spontaneous drift. CLIO retrieval events then act as discrete injections on this flow:

$$C(t_k^+) = \text{CLIO}(C(t_k^-)),$$

so the full system is a hybrid dynamical system: continuous relaxation punctuated by discrete resolution steps. The entropy associated with a synthetic trajectory $\tilde{\gamma}$ now evolves under two competing influences:

$$\frac{d}{dt}S(\tilde{\gamma}) = -\alpha \cdot (\text{exposure rate}) + \beta \cdot (\text{relaxation rate}),$$

where α and β are system-specific constants. Deformation of the constraint surface occurs only when the time-averaged entropy derivative remains negative, meaning the entropy reduction driven by exposure outpaces the entropy inflation driven by relaxation.

This introduces a critical exposure frequency ν^* determined by the ratio β/α . For input sequences with exposure frequency $\nu > \nu^*$, the system remains in the deformation regime of Theorem 4.4(iii): entropy converges to zero and the constraint surface drifts without bound. For $\nu < \nu^*$, relaxation restores entropy between updates faster than exposure reduces it, and the system re-enters the bounded regime of case (ii). The deformation condition is therefore not merely that $S(\tilde{\gamma}_k) \rightarrow 0$ under iteration, but that the exposure cadence must dominate the intrinsic relaxation timescale λ^{-1} .

Within the RSVP framework, this dynamic has a natural interpretation. The entropy field $S(x, t)$ already appears in the state triple as a component of $X(x, t)$. Between CLIO events, it obeys a diffusion-decay equation:

$$\frac{\partial S}{\partial t} = D \nabla^2 S - \lambda S + \text{injection terms},$$

where D governs spatial diffusion of uncertainty across the constraint structure and λ governs decay toward a baseline. CLIO updates act as localized sinks in S , reducing entropy at the position of the synthetic trajectory in field space. The

deformation regime corresponds to sustained localized sinks that outpace diffusion and decay, producing persistent low-entropy basins. The bounded regime corresponds to sinks that are overwhelmed by diffusion and decay between injections, so that no persistent basin forms.

This extension modifies the irreversibility claim of the proof sketch for Theorem 4.4(iii). Deformation is no longer absolutely irreversible but is metastable on timescales set by λ^{-1} . If the perturbation sequence ceases for a time long compared to the relaxation timescale, the system can in principle recover toward the original constraint surface through the continuous flow $-\lambda \nabla V(C)$. Recovery is therefore possible but requires a recovery interval proportional to the degree of deformation achieved and inversely proportional to λ . Social correction and sleep-driven consolidation are natural candidates for processes that increase effective λ , providing the system with mechanisms for partial recovery that the purely discrete model could not represent.

4.6. Minimal Dynamical Characterization of the Deformation Regime

The hybrid system admits a unified characterization in terms of the time evolution of $V(C)$. Let $C(t)$ evolve under continuous relaxation punctuated by discrete CLIO updates at times $\{t_k\}$. The total change in V over an interval $[0, T]$ decomposes as

$$\Delta V = \int_0^T \left. \frac{dV}{dt} \right|_{\text{relax}} dt + \sum_{k: t_k \leq T} [V(C(t_k^+)) - V(C(t_k^-))],$$

where the continuous term is induced by $-\lambda \nabla V(C)$ and the discrete term by CLIO updates. Define the time-averaged drift

$$\bar{V} = \lim_{T \rightarrow \infty} \frac{1}{T} \Delta V.$$

This scalar invariant yields a sharp dynamical classification. When $\bar{V} > 0$, continuous relaxation dominates and trajectories return toward prior basins of V : the recovery regime. When $\bar{V} = 0$, relaxation and injection balance, defining the stability boundary. When $\bar{V} < 0$, discrete updates dominate and the minimizer of V drifts under perturbation: the deformation regime. The critical exposure frequency ν^* is implicitly defined by the condition $\bar{V} = 0$; Theorem 4.4(iii) corresponds to the regime $\nu > \nu^*$ where $\bar{V} < 0$.

Within the RSVP representation, the deformation condition is equivalent to the

existence of a persistent entropy sink:

$$\int_0^T (\text{injection rate} - \lambda S) dt < 0$$

in the long-time limit. Deformation occurs precisely when entropy reduction induced by repeated admissible inputs outpaces both diffusion and intrinsic decay. The system undergoes a phase transition at $\nu = \nu^*$, separating regimes of bounded fluctuation from regimes of cumulative structural drift. This condition is not a heuristic threshold but the sign of a well-defined scalar derived from the Lyapunov structure of Section 4.4.

4.7. Deformation Dynamics

The deformation of Σ is gradual rather than catastrophic. Each accepted synthetic trajectory modifies the admissibility conditions under which future trajectories are evaluated. Formally, $\Sigma_n \rightarrow \Sigma_{n+1} \rightarrow \Sigma_{n+2}$ constitutes a trajectory in constraint space, distinct from any trajectory in field space X . The system's identity — understood as the constraint surface governing reconstruction — is itself a dynamical object subject to continuous modification through reconsolidation.

This has a consequence that is not obvious from static accounts of memory distortion. It is not individual memories that are primarily altered; it is the conditions of admissibility. Deformation acts on the operator \mathcal{A} rather than on any particular γ ; the system's evolution is second-order in the space of admissibility conditions. A synthetic trajectory accepted at step n does not only introduce a false memory; it shifts $\mathcal{A}(C_n)$ so that trajectories similar to $\tilde{\gamma}$ become more readily admitted at step $n + 1$. The deformation is self-propagating: each admitted perturbation lowers the effective obstruction threshold for structurally similar inputs.

4.8. Self-Reinforcing Regime

When an admitted synthetic trajectory modifies the constraint configuration, it becomes part of the structure against which future trajectories are evaluated. If

$$\tilde{\gamma} \in \mathcal{A}(C_n) \Rightarrow \tilde{\gamma} \in \mathcal{A}(C_{n+1}) \text{ with } S_{n+1}(\tilde{\gamma}) < S_n(\tilde{\gamma}),$$

a positive feedback loop is established. Admission reduces entropy on the trace; reduced entropy increases admissibility; increased admissibility facilitates further entropy reduction through subsequent reconsolidation cycles. The trajec-

tory becomes increasingly entrenched in the constraint structure, not because additional evidence supports it but because the structure has been modified to accommodate it.

This is the formal account of what may be called propagandi: constraint surface deformation driven by self-consistent synthetic trajectories. The process does not require deception or explicit manipulation. It requires only that synthetic inputs be locally admissible, repeatable, and self-consistent. This feedback loop operates only within the admissible regime identified in Theorem 4.4; outside this regime, perturbations fail to enter the constraint update and no reinforcement occurs. Once the feedback loop is established within the admissible regime, the deformation is self-sustaining and resistant to correction through CLIO iteration alone, because the modified surface Σ' is itself a stable equilibrium.

4.9. Why Plausibility Dominates Extremity

The theorem and the feedback analysis together explain a pattern that standard accounts of false memory leave underspecified: why implanted false memories tend to be partial, low-confidence, and context-dependent, while naturally acquired confabulations can become robust and identity-constituting.

Extreme synthetic inputs generate high obstruction. They are rejected before entering the perturbation operator and produce no lasting effect on Σ . Plausible synthetic inputs — those whose projection under π closely approximates that of real autobiographical trajectories — generate low obstruction, enter the admitted regime, and initiate entropy reduction. They do not need to replace existing memories; they need only to be compatible with the existing constraint surface. The system is therefore most vulnerable not to contradiction but to near-consistency. A synthetic trajectory that is sixty-five percent structurally identical to a real autobiographical context will meet lower obstruction than a wholly foreign one, be integrated with less disruption, and initiate stronger feedback dynamics than one requiring significant constraint restructuring.

The empirical implication is precise and falsifiable: the model predicts a monotonic relationship between deformation rate and structural proximity under projection π , holding perceptual fidelity constant. Effectiveness is not a function of how dramatically a synthetic input departs from reality but of how closely it approximates the existing autobiographical constraint structure.

4.10. Minimal Formal Structure of Constraint Surface Evolution

We compress the preceding construction into minimal form. Let \mathcal{C} denote the space of admissibility sheaves over X , equipped with the coherence functional $V : \mathcal{C} \rightarrow \mathbb{R}_{\geq 0}$, $V(C) = \|\text{Obs}\|^2 + S$. The entropy term is understood as a functional derivative of V restricted to the synthetic trajectory: $S(\tilde{\gamma}) = \delta V / \delta \tilde{\gamma}$, so that V generates both the constraint evolution and the entropy dynamics on trajectories within a single potential. CLIO induces a discrete descent operator $\mathcal{T}(C) = \text{CLIO}(C)$ satisfying $V(\mathcal{T}(C)) \leq V(C)$, with equality at fixed points. The constraint surface evolves as $C_{n+1} = \mathcal{T}(C_n)$, $\Sigma_n = \Gamma(X, \mathcal{F}_{C_n})$.

A synthetic perturbation $\tilde{\gamma}$ induces a perturbed operator $\mathcal{T}_{\tilde{\gamma}}(C) = \text{CLIO}(C \cup \tilde{\gamma})$ and a deformation operator on surfaces $\mathcal{D}_{\tilde{\gamma}} : \Sigma(C) \mapsto \Sigma(\mathcal{T}_{\tilde{\gamma}}(C))$. Identity deformation corresponds to a trajectory in \mathcal{C} under iterated perturbed descent: $C_n \mapsto \mathcal{T}_{\tilde{\gamma}_n} \circ \dots \circ \mathcal{T}_{\tilde{\gamma}_1}(C_0)$. Identity is therefore not a fixed point in \mathcal{C} but the minimizer of a functional whose landscape is itself modified by admissible inputs. Deformation is motion of this minimizer under perturbation, not accumulation of states within a fixed admissible set.

Recovery from deformation is not symmetric with its onset. During deformation, the system is driven along a sequence of moving minima in V . During recovery, it descends toward a prior basin that may no longer be the global minimum after deformation. If the deformation has shifted the system across a barrier in V , relaxation alone cannot restore the original surface: basin accessibility, not merely decay rate, determines whether recovery is possible. Metastability can therefore become effectively permanent without additional perturbation, and the recovery timescale is controlled jointly by λ^{-1} and the height of the barrier separating the deformed configuration from the original basin.

5. Synthetic Reification and the Epistemic Vulnerability of Unverifiable Content

The argument of Section 4 establishes the conditions under which constraint surface deformation must occur: a synthetic trajectory must generate near-zero obstruction in $\check{H}^1(\mathcal{U}, \mathcal{F}_n)$, drive entropy $S(\tilde{\gamma} \mid C_n)$ toward zero through repetition, and do so within the admissible regime. The present section identifies why synthetic reification of unverifiable content satisfies these conditions with particular force, and why the content domains where humans are most dependent on mediated representation — news events, scientific findings, historical

sequences, social norms, tool demonstrations, physical phenomena — are precisely those where the cognitive system has the least independent constraint to resist deformation. Each relevant property of high-fidelity synthetic content corresponds directly to one of the formal variables governing admissibility.

The mechanism generalizes the availability heuristic into a structural claim. Tversky and Kahneman identified that ease of mental simulation inflates perceived probability and importance. The present framework specifies why: vivid, concrete representations reduce entropy $S(\tilde{\gamma})$ by narrowing the space of compatible reconstructions, and they reduce obstruction $\text{Obs}(\tilde{\gamma}, C_n)$ by populating constraint patches with well-formed local sections. The heuristic is not a bias layered on top of otherwise sound reasoning; it is the direct output of coherence-based acceptance applied to high-density inputs. Reification — the rendering of abstract or unverifiable claims in perceptually concrete form — is not merely persuasive rhetoric but a mechanism for satisfying the formal conditions of Theorem 4.4(iii).

5.1. Perceptual Density and Entropy Reduction

The first condition for entry into the deformation regime is low entropy $S(\tilde{\gamma})$. High entropy corresponds to uncertainty in reconstruction, leading to oscillatory CLIO dynamics and eventual rejection. Low entropy corresponds to convergence under a negative uncertainty gradient and stable integration. Generative video systems produce outputs with unusually high perceptual density: fine-grained visual detail, continuous motion, consistent lighting and depth, and synchronized audio.

This perceptual density functions as a direct reduction in entropy. A synthetic trajectory $\tilde{\gamma}$ generated with sufficient fidelity presents a densely specified candidate state in model space, reducing the degrees of freedom available during reconstruction. Formally, the space of compatible trajectories under $\pi^{-1}(m)$ is narrowed, even though π^{-1} remains non-invertible, yielding

$$S(\tilde{\gamma}) \ll S(\gamma_{\text{imagined}})$$

for comparable imagined or verbally described scenarios. The CLIO resolution loop therefore encounters a low-entropy candidate that satisfies convergence conditions without requiring additional constraint resolution. This directly places such inputs in the admitted regime of Theorem 4.4 and, under repetition, in the deformation regime where $S(\tilde{\gamma}) \rightarrow 0$.

Imagination and prior synthetic media — text, still images, low-fidelity video — produce sparse local sections that fail to satisfy multiple patch constraints simultaneously. The resulting cocycles are large, and the cohomology class $[\omega]$ is nontrivial. This is the technical reason that normal imagination does not typically produce memory-feeling: it lacks the density required to generate near-zero obstruction across the full constraint cover.

5.2. Structural Proximity and the Unverifiable Domain

The second condition for deformation is low obstruction: $\text{Obs}(\tilde{\gamma}, C_n) \approx 0$. Obstruction is minimized when the synthetic input projects close to existing structures in the constraint surface — when it is not merely coherent in isolation but coherent with what the subject already believes, knows, or has been taught. This is the specific vulnerability of content about things one cannot directly verify.

For domains accessible to direct experience — the layout of a familiar room, the texture of a surface, the behavior of a known person — the subject has dense prior constraint that resists synthetic inputs inconsistent with it. The cocycles ω_{ij} on overlaps between personal experience patches and the synthetic input will be large if the input conflicts with established structure, generating obstruction that drives rejection. But for content no one can directly experience — the interior of a particle accelerator, the dynamics of a historical battle, the mechanisms of a disease process, the social norms of an unfamiliar community — the constraint surface has sparse coverage. There is no dense prior to generate conflicting cocycles. Synthetic reifications of such content encounter near-zero obstruction by default:

$$\omega_{ij} \approx 0 \quad \text{across } \mathfrak{U}$$

because the existing patches $\mathcal{F}_n(U_i)$ contain little material to conflict with. The cohomology class $[\omega]$ approaches triviality not because the synthetic input is particularly well-crafted but because the constraint surface is structurally sparse in the relevant domain. This is the formal basis for why synthetic reification of news, scientific findings, and historical events is more effective than synthetic reification of personally verifiable claims: the former faces near-zero resistance, the latter faces dense prior constraint.

5.3. Repeatability and Entropy Reduction

The third and decisive condition is convergence $S(\tilde{\gamma}_k) \rightarrow 0$ through repeated exposure. Each exposure constitutes one CLIO cycle:

$$C_{n+1} = \text{CLIO}(C_n \cup \tilde{\gamma}).$$

Repeated application yields a sequence of updates in which the entropy associated with $\tilde{\gamma}$ decreases monotonically:

$$S_{n+1}(\tilde{\gamma}) < S_n(\tilde{\gamma}).$$

Generative video systems support exact replay: a given synthetic sequence can be viewed repeatedly without degradation, variation, or contextual drift. Unlike genuine episodic memories, which degrade and vary across recalls due to reconstruction noise, the synthetic input is stable. Each replay drives $S(\tilde{\gamma})$ further toward zero without the entropy inflation that genuine reconsolidation introduces. Because each individual perturbation contributes negligibly to the coherence functional $V(C)$ while cumulatively displacing its minimizer, the system enters the regime of continuous deformation. Without repeatability, perturbations remain in the bounded regime; with repeatability, they converge into the deformation regime.

5.4. Domain Conditioning and Narrative Consistency

The fourth feature determining how rapidly deformation conditions are satisfied is domain-specific conditioning: synthetic content that is calibrated to existing beliefs, prior exposures, and the narrative structures of a given knowledge domain. Let d_π denote a metric on model space induced by π . Content conditioned on what the subject already believes about a domain yields

$$d_\pi(\pi(\tilde{\gamma}), \pi(\gamma)) \rightarrow 0$$

for trajectories γ drawn from existing constraint structures. Synthetic news that is consistent with the subject's prior political beliefs, synthetic scientific content that fits existing conceptual frameworks, synthetic historical narratives that cohere with established timeline structures: each reduces obstruction by aligning with existing constraint patches rather than conflicting with them. The system is biased toward generating perturbations already near the interior of the admissible region, directly satisfying the condition of Section 4.7.

This explains why the most epistemically dangerous synthetic content is not fabrication inconsistent with established facts — which faces high obstruction and is often rejected — but fabrication that is near-consistent with them: alternate histories that differ from documented ones in plausible ways, scientific findings that extend accepted frameworks in coherent directions, social norms that are adjacent to familiar ones. Near-consistency is the condition for low obstruction; low obstruction is the condition for deformation.

5.5. Developmental Vulnerability and the Formation of Epistemic Norms

The analysis above applies to adult subjects with established constraint surfaces. The concern is substantially amplified in developmental contexts, where the constraint surface itself is being formed.

A child or novice encountering a domain for the first time has no prior dense constraint against which to evaluate synthetic reifications. The admissibility sheaf \mathcal{F}_n is sparse: few patches $\mathcal{F}_n(U_i)$ contain material sufficient to generate conflicting cocycles. In this condition, virtually any coherent, perceptually dense synthetic input satisfies $\text{Obs}(\tilde{\gamma}, C_n) \approx 0$ and enters the deformation regime immediately. The constraint surface does not resist the synthetic input; it is formed partly by it.

This has a structural consequence that goes beyond any individual false belief. The admissibility conditions $\mathcal{A}(C_n)$ that govern future evaluations are themselves being established during the period of exposure. Synthetic inputs accepted early in development do not merely introduce specific false trajectories; they shift the conditions under which future inputs will be evaluated. A constraint surface formed in part by synthetic reifications of scientific phenomena, social norms, or historical events will have different admissibility conditions than one formed through verified experience and reliable testimony. The deformation operator $\mathcal{D}_{\tilde{\gamma}}$ acts not on a mature surface but on the generative process that produces it.

The developmental concern also applies to the formation of trust itself. The calibration of which sources, institutions, and representational formats are reliable is itself a learning process that proceeds through exposure. If synthetic content occupying the low-obstruction, low-entropy regime is present throughout that calibration process, the resulting trust structure will be shaped by it. This is not a side effect of synthetic media but a direct consequence of the architecture: any system that resolves for coherence rather than provenance will weight reliable

and unreliable sources similarly when both satisfy the coherence constraints, and the subject cannot distinguish them by introspection alone.

5.6. Spatial Coherence and Cognitive Map Recruitment

A further feature of generative video systems that standard synthetic media do not share is the generation of navigable or quasi-navigable spatial environments. Scenes exhibit consistent geometry, depth cues, and viewpoint continuity across frames, aligning with the cognitive map structures described by O’Keefe and Nadel [12], which encode environments as relational spatial configurations rather than static images.

Within the RSVP framework, spatial coherence corresponds to alignment in the vector field \mathbf{v} governing trajectory structure. A synthetic trajectory that preserves directional consistency and spatial relations produces a candidate $\tilde{\gamma}$ whose vector components match those of real trajectories:

$$\mathbf{v}_{\tilde{\gamma}} \approx \mathbf{v}_{\gamma}.$$

This alignment reduces both entropy and obstruction simultaneously, as fewer adjustments are required during reconstruction to satisfy directional constraints. Episodic memory is organized in part around spatial context [13]; a synthetic input that presents coherent spatial structure engages the hippocampal system in the same mode as genuine environmental experience rather than as a static perceptual stimulus. The resulting trajectory is processed as a coherent path through state space, further increasing its admissibility and stability under CLIO iteration. The distinction between viewing a scene and implicitly inhabiting a space corresponds, formally, to the difference between engaging a subset of constraint patches and engaging the full cover \mathcal{U} .

5.7. Narrative Coherence and Higher Obstruction Classes

The conditions established in subsections 5.1 through 5.5 address the primary obstruction $[\omega] \in \check{H}^1(\mathcal{U}, \mathcal{F}_n)$. Higher cohomology classes $[\omega] \in \check{H}^p$ for $p \geq 2$ capture more subtle inconsistencies that only manifest when three or more constraint patches are considered simultaneously: temporal loops, causal contradictions across episodes, or violations of narrative coherence that survive pairwise checking but fail under triple or higher intersection conditions.

Contemporary generative systems that condition on user prompts, memory em-

beddings, or ongoing conversational history produce synthetic trajectories that respect not only local patches but also higher-order consistency on triple and higher intersections. This systematically drives higher cochain norms toward zero, rendering primary obstruction classes killable at low entropy cost even when secondary obstructions are present. The result is that the full gluing condition $[\omega] \approx 0$ in $\check{H}^*(\mathcal{U}, \mathcal{F}_n)$ can be approximately satisfied across all cohomological degrees simultaneously, not only at the level of pairwise overlaps.

The significance of this is that social verification, which operates in part by probing triple and higher consistency — does this story cohere with what you told me before, and what we both remember from that time — loses much of its corrective power against well-conditioned synthetic inputs. The synthetic trajectory has already satisfied the consistency conditions that social cross-referencing would expose. This connects directly to the argument of Section 6, where social correction is formalized as external constraint enforcement operating on the shared sheaf.

5.8. Generative Systems as Systematic Deformation Drivers

Because synthetic reification of unverifiable content jointly satisfies low obstruction (from sparse prior constraint in the relevant domain), low entropy (from perceptual density and narrative consistency), and convergent entropy reduction (from repeatability and domain conditioning), it systematically produces inputs lying in the third regime of Theorem 4.4. The roles of entropy and obstruction are distinct. Perceptual density primarily reduces entropy $S(\tilde{\gamma})$ by narrowing the space of compatible reconstructions. Domain proximity primarily reduces obstruction $\text{Obs}(\tilde{\gamma}, C_n)$ by aligning with existing constraint patches. Repeatability drives entropy toward zero through iterated CLIO cycles. Narrative conditioning reduces the magnitude of higher cohomology classes. No single property is sufficient alone, but content optimized for credibility, engagement, and domain relevance tends to satisfy all conditions simultaneously. This is not a deliberate exploit of the cognitive architecture but a structural consequence of optimizing for epistemic persuasiveness.

The framework therefore predicts selective, gradual, and self-reinforcing integration of synthetic autobiographical content, with deformation rate governed by the frequency and quality of exposure and by the degree of alignment between the synthetic input and the subject's existing constraint surface. This completes the deductive bridge: Sections 2 and 3 establish non-invertibility and the

absence of provenance checking; Section 4 derives the precise conditions under which deformation must occur; Section 5 demonstrates that current generative systems are engineered, whether intentionally or by optimization pressure, to meet those conditions at scale.

5.9. Implications for Empirical Testing

The formal correspondence between system properties and deformation conditions yields precise empirical predictions. The rate and magnitude of constraint surface deformation should correlate with structural proximity under π and with the rate of entropy reduction induced by repeated exposure. Entropy reduction is operationalizable as decreased reconstruction latency and increased confidence ratings across successive exposures to the same synthetic input; obstruction is operationalizable as cross-context inconsistency detection rates, measurable by probing whether subjects notice contradictions when the synthetic trajectory is embedded in narratively varied contexts. Holding perceptual fidelity constant, synthetic inputs that more closely approximate existing autobiographical trajectories should produce greater deformation than those that diverge more dramatically. Personalized inputs should produce deformation at lower exposure thresholds, with higher stability, and with greater resistance to social correction than generic synthetic inputs; this implies deformation rate should be a measurable function of personalization degree, holding perceptual density and exposure frequency constant. Conversely, limiting repeatability or increasing variability across exposures should maintain perturbations within the bounded regime of Theorem 4.4(ii), which predicts recovery of the original constraint configuration after cessation of the perturbation sequence. These predictions follow directly from the stability theorem and are independent of the specific implementation details of any generative system.

6. Social Correction as External Constraint Enforcement

The preceding section established that contemporary generative systems systematically produce synthetic trajectories satisfying the conditions of Theorem 4.4(iii), driving entropy toward zero and minimizing obstruction within a single-agent constraint configuration. The result is unbounded deformation of the autobiographical constraint surface under iterated reconsolidation. The present section introduces the countervailing mechanism: social interaction as external constraint enforcement. Where generative systems operate within a single constraint surface to collapse entropy and obstruction, social processes couple mul-

tuple such surfaces, reintroducing both through cross-agent consistency requirements.

Let $C_n^{(i)}$ and $C_n^{(j)}$ denote the constraint configurations of two agents i and j , with corresponding admissibility sheaves $\mathcal{F}_n^{(i)}$ and $\mathcal{F}_n^{(j)}$. Social interaction introduces a coupling operator

$$\mathcal{C}_{ij} : (C_n^{(i)}, C_n^{(j)}) \longmapsto (C_{n+1}^{(i)}, C_{n+1}^{(j)}),$$

which enforces partial alignment between the two constraint configurations. At the level of sheaves, this corresponds to imposing compatibility conditions on overlaps of autobiographical domains: shared events, common environments, or mutually referenced narratives. Formally, the coupled system induces a fiber product of sheaves over shared subdomains,

$$\mathcal{F}_n^{(ij)} = \mathcal{F}_n^{(i)} \times_{\mathcal{F}_n^{(\text{shared})}} \mathcal{F}_n^{(j)},$$

where $\mathcal{F}_n^{(\text{shared})}$ encodes constraints derived from jointly accessible or mutually verifiable experience. Under this coupling, admissibility is no longer evaluated within a single constraint surface but relative to the joint system. A trajectory $\tilde{\gamma}$ admitted in $C_n^{(i)}$ must satisfy not only $\tilde{\gamma} \in \mathcal{A}(C_n^{(i)})$ but also compatibility with $\mathcal{F}_n^{(j)}$ on shared domains.

The obstruction functional therefore acquires an inter-agent component:

$$\text{Obs}_{ij}(\tilde{\gamma}) = \text{Obs}(\tilde{\gamma}, C_n^{(i)}) + \text{Obs}(\tilde{\gamma}, C_n^{(j)}),$$

together with higher-order terms arising from mismatches across the fiber product. Even if $\text{Obs}(\tilde{\gamma}, C_n^{(i)}) \approx 0$, incompatibility with $C_n^{(j)}$ may yield $\text{Obs}_{ij}(\tilde{\gamma}) \gg 0$, preventing integration into the coupled system without additional constraint restructuring. This mechanism lifts the obstruction problem to a higher-dimensional constraint space: within a single agent, obstruction is measured relative to $\mathcal{F}_n^{(i)}$ alone; under coupling, the relevant cohomology classes are defined over the joint cover induced by $\mathfrak{U}^{(i)} \cup \mathfrak{U}^{(j)}$. Synthetic trajectories that are locally admissible within one cover may fail to glue across the combined cover, generating nontrivial cohomology classes that cannot be killed without increasing entropy. Social correction therefore reintroduces both obstruction and entropy into the reconstruction process, pushing perturbations from regime (iii) of Theorem 4.4 back toward regimes (ii) or (i).

The effect on entropy is equally direct. Within a single-agent system, repeated exposure to a synthetic trajectory produces monotonic entropy reduction, $S_{n+1}(\tilde{\gamma}) < S_n(\tilde{\gamma})$, driving the system toward convergence. Under social coupling, disagreement between agents introduces uncertainty into the reconstruction process. A trajectory stable within $C_n^{(i)}$ may be challenged by $C_n^{(j)}$, generating alternative candidate reconstructions and increasing the effective entropy of the state. The entropy functional becomes

$$S_{ij}(\tilde{\gamma}) = S^{(i)}(\tilde{\gamma}) + S^{(j)}(\tilde{\gamma}) + S^{(\text{interaction})}(\tilde{\gamma}),$$

where the interaction term captures uncertainty induced by cross-agent inconsistency. This term is strictly positive when constraint configurations differ and counteracts the entropy reduction produced by repetition. The CLIO dynamics of the coupled system therefore no longer exhibit monotonic convergence for the synthetic trajectory; instead, they return to oscillatory or unstable regimes characteristic of failed reconsolidation.

The stabilizing effect of social interaction can be understood as an externalization of the Lyapunov functional introduced in Section 4. The coherence functional $V(C)$ is no longer minimized over a single configuration space but over a coupled space subject to additional constraints. Synthetic perturbations that would shift the minimum of V in the single-agent case encounter a modified landscape in the coupled system, with additional curvature induced by inter-agent consistency requirements. This increased curvature raises the effective cost of displacement, counteracting the cumulative drift described in Theorem 4.4(iii).

This stabilizing mechanism has, however, a structural limitation. If multiple agents are exposed to the same synthetic trajectory and update their constraint configurations in parallel, the coupling operator no longer introduces independent constraints. Instead, the systems evolve toward a shared but deformed constraint surface. Formally, if

$$C_n^{(i)} \approx C_n^{(j)} \quad \text{and} \quad \tilde{\gamma} \in \mathcal{A}(C_n^{(i)}) \cap \mathcal{A}(C_n^{(j)}),$$

then the joint obstruction $\text{Obs}_{ij}(\tilde{\gamma})$ remains small, and the perturbation is admitted in both systems. Repetition across agents drives entropy toward zero in each configuration simultaneously, and the coupled system enters the same deformation regime as the single-agent case. Social coupling in this context produces correlated drift rather than correction. The shared synthetic environment does not introduce the independent constraint variance that makes coupling cor-

rective; it removes it.

The distinction between independent and correlated constraint configurations is therefore the critical variable. We define a corrigibility condition as $\text{Var}(C^{(i)}, C^{(j)}) > 0$: the constraint configurations of interacting agents must exhibit nonzero variance for coupling to function as correction. When this condition holds, disagreement between agents introduces independent obstruction and entropy inflation, counteracting the convergence dynamics of regime (iii). Social correction is effective only when agents provide partially independent constraint surfaces, such that inconsistencies in one configuration generate obstruction in another. When $\text{Var}(C^{(i)}, C^{(j)}) \rightarrow 0$, the system enters the synchronization regime: the corrective mechanism is replaced by a synchronization process in which multiple constraint surfaces deform in parallel toward a common equilibrium Σ' . Social correction is therefore not binary but a function of constraint variance across interacting agents. This predicts that shared synthetic environments do not mitigate deformation but amplify it by increasing exposure frequency and eliminating the cross-agent variance that makes coupling corrective.

The empirical implication is that resistance to synthetic trajectory integration should correlate with the diversity and independence of social interactions. Systems in which agents encounter heterogeneous inputs and maintain partially distinct constraint configurations will exhibit higher effective obstruction and entropy, limiting deformation. Systems characterized by homogeneous exposure and synchronized conditioning will exhibit reduced obstruction and accelerated drift. The measurable proxy for constraint independence is cross-agent inconsistency detection: if agents exposed to the same synthetic input consistently agree rather than challenge each other's reconstructions, the coupling operator has become a synchronization mechanism rather than a correction mechanism, and the system has entered the amplifying regime.

In summary, social interaction operates as external constraint enforcement by coupling individual admissibility sheaves and lifting the obstruction problem to a multi-agent domain. This coupling reintroduces entropy and increases obstruction, counteracting the convergence dynamics that drive deformation in the single-agent case. The effectiveness of this mechanism depends entirely on the independence of constraint configurations across agents. When that independence is preserved, social interaction is the primary structural defense against constraint surface deformation. When it is lost — through shared exposure, synchronized conditioning, or homogeneous input environments — social

processes cease to function as correction and instead become a channel for coordinated deformation.

7. Limits, Falsifiability, and Experimental Program

The preceding sections establish a formal equivalence between memory reconsolidation and CLIO-style recursive resolution, derive a stability theorem governing constraint surface deformation under synthetic perturbation, and identify contemporary generative systems as inputs that satisfy the deformation regime. The strength of the framework lies in the fact that its central claims are not merely descriptive but structurally necessary given the assumptions of non-invertible projection and coherence-based resolution. This section isolates those assumptions, delineates the limits of the model, and formulates explicit empirical tests capable of falsifying its core predictions.

7.1. Minimal Assumptions

The deformation result depends on three minimal structural assumptions. First, projection from field space to model space is non-invertible, so that multiple trajectories correspond to indistinguishable representations under π . Second, reconstruction proceeds by selection over admissible trajectories rather than retrieval of a unique stored path. Third, the acceptance criterion is coherence-based, with no direct dependence on provenance. If any of these assumptions fails, the deformation mechanism is weakened or eliminated. If π were invertible, provenance would be recoverable and synthetic trajectories could be rejected on structural grounds. If reconstruction were reproductive rather than generative, synthetic inputs would not enter the reconstruction process as candidate trajectories. If acceptance depended explicitly on provenance, coherence alone would not suffice for integration. The model therefore makes a strong conditional claim: given these three properties, deformation under low-obstruction, low-entropy perturbation is structurally unavoidable.

7.2. Limits of the Formalization

The RSVP representation abstracts away from biological implementation details, treating the system at the level of constraint dynamics rather than neural circuitry. This abstraction allows the argument to apply across a wide class of systems exhibiting reconstruction under compression, but it does not specify the exact neural substrates of the fields (Φ, \mathbf{v}, S) or the operator CLIO. The the-

ory therefore does not predict micro-level mechanisms such as synaptic changes or specific regional activation patterns beyond those already associated with reconsolidation.

A second limitation concerns the topology of the constraint space. The deformation analysis assumes that the space of admissibility sheaves admits a sufficiently smooth moduli structure to support continuous drift under small perturbations. Pathological cases in which the moduli space is highly disconnected or exhibits large discrete jumps would violate the continuous deformation regime described in Theorem 4.4.

A third limitation concerns temporal scale. The analysis treats reconsolidation as an iterated process indexed by discrete updates but does not specify the timescale over which entropy reduction and constraint updates occur. In biological systems, reconsolidation is gated by retrieval and may be subject to consolidation windows, sleep cycles, and neuromodulatory states. These factors may slow, gate, or partially reverse deformation without altering its structural possibility.

7.3. Falsifiability Conditions

The theory makes three primary falsifiable predictions, each corresponding to one component of the stability theorem. First, deformation rate must correlate with structural proximity under projection π : holding perceptual fidelity constant, synthetic inputs that more closely match a subject's autobiographical structure should produce larger shifts in subsequent admissibility judgments. Failure to observe this monotonic relationship would falsify the central role of obstruction as defined by $\text{Obs}(\tilde{\gamma}, C_n)$. Second, repeatability must drive entropy reduction and convergence: repeated exposure to an identical synthetic trajectory should decrease reconstruction latency, increase subjective confidence, and increase cross-context consistency of recall. If repeated exposure does not produce these effects, the identification of CLIO dynamics with reconsolidation would be undermined. Third, social independence must determine resistance to deformation: subjects embedded in heterogeneous social environments should exhibit higher rates of inconsistency detection and lower rates of synthetic integration than subjects exposed to homogeneous or synchronized inputs. If social coupling does not increase effective obstruction or entropy under disagreement, the model of external constraint enforcement fails. Each of these predictions is independent; failure of any one would require revision of the corresponding

component without necessarily invalidating the entire framework.

7.4. Experimental Program

These falsifiability conditions translate directly into an experimental program. A first class of experiments tests projection proximity by exposing subjects to synthetic autobiographical scenarios varying systematically in structural similarity to their real experiences while holding perceptual fidelity constant. Structural similarity is manipulated through self-representation, environmental familiarity, and narrative continuity, with integration measured by subsequent recall, confidence ratings, and incorporation into future reconstructions. A second class tests entropy dynamics under repetition by presenting identical synthetic sequences across multiple exposures and measuring recall latency, confidence, neural pattern similarity across retrievals, and cross-context consistency. The model predicts monotonic entropy decrease observable as increased stability and reduced variability across reconstructions. A third class tests social coupling by assigning subjects to groups with either independent or shared exposure to synthetic inputs, then measuring inconsistency detection, confidence, and memory integration after interaction. A fourth class combines these manipulations to test the transition between bounded deformation under Theorem 4.4(ii) and unbounded drift under Theorem 4.4(iii): the critical prediction is that the same input can fall into different stability regimes depending on social coupling conditions alone.

A fifth class of experiments directly tests the temporal relaxation dynamics. The model predicts that massed and spaced exposure to identical synthetic inputs should produce qualitatively different deformation regimes, not merely quantitative differences in integration strength. At exposure frequencies below ν^* , relaxation between presentations restores entropy faster than successive exposures reduce it, maintaining perturbations within the bounded regime and allowing recovery if exposure ceases. At frequencies above ν^* , the system crosses into regime (iii) and undergoes cumulative deformation that persists after exposure ends. This predicts a phase-transition signature in the data: an abrupt change in integration probability as a function of inter-stimulus interval. Critically, the model predicts not only a change in mean confidence or latency at this threshold but a qualitative reversal in variance: below ν^* , the variance of reconstructions should increase over time as relaxation introduces dispersion; above ν^* , variance should collapse as entropy converges toward zero and reconstructions stabilize. This variance flip is a distinctive prediction that smooth

monotonic models of repetition cannot produce and that is directly observable in behavioral data without requiring neural measurement.

8. Conclusion

The analysis developed in this paper can be compressed into a single structural claim. Let a system satisfy three conditions: projection from state space to representation space is non-invertible, reconstruction proceeds by selection over admissible trajectories rather than retrieval of a fixed trace, and acceptance is governed by coherence rather than provenance. Then any sequence of inputs that simultaneously minimizes obstruction under the induced admissibility sheaf and drives the associated entropy functional toward zero under iteration will not merely be incorporated into reconstruction but will deform the constraint surface that defines admissibility itself. This claim is independent of the specific domain in which the system is instantiated. It applies to belief formation about scientific findings, historical events, social norms, and technical phenomena, as well as to autobiographical memory, insofar as these processes share the same structural properties.

The RSVP formalization provides one representation of this class of systems; the CLIO mapping identifies the corresponding computational dynamics; Theorem 4.4 establishes the conditions under which deformation is bounded or unbounded. Synthetic reification of unverifiable content satisfies these conditions not accidentally but as a consequence of the same properties that make it epistemically persuasive: perceptual density, narrative consistency, domain coherence, and repeatability. The content that is most effective for learning and communication is, by the same mechanisms, most effective for constraint surface deformation.

The resulting perspective reframes the problem of synthetic influence on belief. The relevant question is not whether a given representation is true or false, nor whether a subject can in principle identify it as synthetic, but whether it lies within the low-obstruction, low-entropy region of the admissible set. What the framework makes explicit is that truth is not a fixed point of the dynamics — coherence is. The reconstruction operator \mathcal{R} converges on states that are maximally coherent, not maximally accurate; the Lyapunov functional V is minimized by constraint surfaces that are internally consistent, not externally verified. Synthetic reifications that achieve near-zero obstruction and entropy are therefore not exploiting a weakness in the cognitive system but its fundamental

operating principle.

The developmental consequence follows directly. Children and domain novices encounter synthetic reifications before they have established dense prior constraint. Their constraint surfaces are formed partly by the inputs they receive during formation. A cognitive ecology that includes high-fidelity synthetic reifications of unverifiable content throughout the developmental period does not merely introduce specific false beliefs; it shapes the admissibility conditions under which all future inputs will be evaluated. Under the formulation developed here, this is not an anomaly but a theorem: given the stated assumptions, it is the expected behavior of any system that reconstructs under compression and resolves under coherence. The limits of this conclusion coincide exactly with the limits of those assumptions, and its empirical status is determined by the falsifiability conditions outlined in the preceding section.

References

- [1] Adam Zeman, Michaela Dewar, and Sergio Della Sala, “Lives without imagery — congenital aphantasia,” *Cortex*, 73, 2015, pp. 378–380.
- [2] Endel Tulving, *Elements of Episodic Memory*. Oxford University Press, 1983.
- [3] Gilles Fauconnier and Mark Turner, *The Way We Think: Conceptual Blending and the Mind’s Hidden Complexities*. Basic Books, 2002.
- [4] David E. Meyer and Roger W. Schvaneveldt, “Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations,” *Journal of Experimental Psychology*, 90(2), 1971, pp. 227–234.
- [5] Newman Cheng, Gordon Broadbent, and William Chappell, “Cognitive Loop via In-Situ Optimization: Self-Adaptive Reasoning for Science,” *arXiv preprint arXiv:2508.02789*, 2025. DOI: 10.48550/arXiv.2508.02789.
- [6] Elizabeth F. Loftus and Jacqueline E. Pickrell, “The Formation of False Memories,” *Psychiatric Annals*, 25(12), 1995, pp. 720–725.
- [7] Marcia K. Johnson, Shahin Hashtroudi, and D. Stephen Lindsay, “Source Monitoring,” *Psychological Bulletin*, 114(1), 1993, pp. 3–28.
- [8] Daniel L. Schacter, *Searching for Memory: The Brain, the Mind, and the Past*. Basic Books, 1996.

- [9] Karim Nader, Glenn E. Schafe, and Joseph E. LeDoux, "Fear memories require protein synthesis in the amygdala for reconsolidation after retrieval," *Nature*, 406, 2000, pp. 722–726.
- [10] Oliver Hardt, Einar Ör Einarsson, and Karim Nader, "A bridge over troubled water: Reconsolidation as a link between cognitive and neuroscientific memory research traditions," *Annual Review of Psychology*, 61, 2010, pp. 141–167.
- [11] David B. Pillemer, *Momentous Events, Vivid Memories*. Harvard University Press, 1998.
- [12] John O'Keefe and Lynn Nadel, *The Hippocampus as a Cognitive Map*. Oxford University Press, 1978.
- [13] Jeffrey M. Zacks et al., "Event perception: A mind-brain perspective," *Psychological Bulletin*, 133(2), 2007, pp. 273–293.