Reconstruction Stability and Synthetic Contamination

Memory, Identity, and Constraint Closure
Under Adversarial Observation Injection

Flyxion

Independent Researcher

April 2026

Abstract. Human memory is not retrieval from a stable archive but iterative reconstruction under partial constraints. This essay formalizes that observation and traces its consequences across a hierarchy of increasingly adversarial conditions. We introduce autobiographical reconstruction as a constrained fixed-point problem, define the notion of synthetic projection contamination, and prove a theorem establishing the conditions under which high-fidelity synthetic observations can induce stable but non-realizable memory states. We then develop a connected series of results: the decomposition of the notification index for collaborative systems, the residual energy functional governing memory persistence, an interrupted closure lemma formalizing the Zeigarnik effect, a proposition on closure without grounding that unifies dream false memory, collective false memory, and large language model hallucination, and a synchronization principle governing constraint drift across decoupled instances. Throughout, formal results are grounded in observations drawn from cognitive science, cultural artifacts, everyday encoding practices, and the product-level dynamics of synthetic media platforms. The central claim is that truth in a reconstruction-based system is not correspondence to a stored record but survival as a fixed point under a sufficiently rich, robust, and externalized constraint closure process.

Introduction: Memory as Reconstruction

The standard picture of memory is archival. An experience is encoded, stored in some neural substrate, and later retrieved, approximately intact. On this picture, false memory is a malfunction: the retrieval mechanism has returned the wrong file, or the file itself has been corrupted. The task of memory research is to characterize these failure modes and explain why the archive is less reliable than it appears.

This essay proceeds from a different starting assumption. Memory is not retrieval from an archive but a process of iterative reconstruction under partial constraints. Each act of remembering is a new synthesis, constrained by prior encodings, sensory cues, semantic expectations, and the current state of the agent. There is no privileged original to which recall can be compared. There is only the reconstruction process, and the fixed points it happens to reach.

This framing has significant consequences. It means that the stability of a memory is not a function of storage fidelity but of constraint richness. Memories that are embedded within dense networks of mutual constraints—semantic, social, perceptual, procedural—are difficult to deform because deforming them would violate many constraints simultaneously. Memories that rest on sparse or weak constraint sets are malleable, susceptible to suggestion, reconsolidation, and drift. The question "how accurate is this memory?" becomes "under what constraint system was this reconstruction produced?"

The contemporary relevance of this shift is acute. High-fidelity generative systems can now produce synthetic media—video, audio, and image—that satisfies the brain's projection-level admissibility criteria. Such systems produce signals that are rich, self-referential, and replayable, activating the same reconstruction pathways as genuine experience while lacking any correspondence to an actual event. Understanding the danger these systems pose requires not a theory of storage corruption, but a theory of constraint contamination. That is what this essay attempts to provide.

We proceed as follows. We begin with the formal setup of autobiographical reconstruction as a constrained fixed-point problem and define the key notion of synthetic projection contamination. We prove the central Reconstruction Stability Theorem, then

develop a series of connected results covering provenance, collective memory, encoding modality, residual energy, and constraint synchronization. Throughout, we situate the formal developments within specific phenomena: false memories induced by film, shared hallucinations in social settings, the encoding advantages of formal notation, and the drift of personalized input systems after centralized updates cease.

Autobiographical Reconstruction as a Constrained Fixed Point

Let the internal autobiographical state be a tuple $X \in \mathcal{A}$, where $\mathcal{A}$ is the space of admissible states equipped with appropriate regularity conditions. In the RSVP framework, this state has components corresponding to scalar, vector, and energetic fields over the agent's experienced domain, but for present purposes the precise structure of $\mathcal{A}$ is less important than the existence of a well-posed constraint system acting on it.

Let the brain's reconstruction process be modeled as a consistency operator

$$C : \mathcal{A} \to \mathcal{A}, \tag{1}$$

such that a coherent autobiographical state satisfies $X^* = C(X^*)$. The fixed points of C are the stable memories: the states that the system, under its current constraint set, can maintain without internal contradiction.

Observations are given by a projection family $\{\Pi_i\}_{i \in I}$, where each $\Pi_i : \mathcal{A} \to \mathcal{Y}_i$ extracts a particular aspect of the state. The observed data $\{y_i\}$ may include perceptual input, emotional response, spatial context, social signals, and internal cues. Reconstruction corresponds to solving the variational problem

$$X^* = \arg\min_{X \in \mathcal{A}} \left(\sum_{i \in I} \|\Pi_i(X) - y_i\|^2 + \mathcal{R}(X) \right), \tag{2}$$

where $\mathcal{R}$ encodes narrative coherence, physical plausibility, identity continuity, and other structural regularities. The fixed point X^* is the state that best reconciles all available

observations under the system's internal coherence norms.

The crucial assumption of the classical framework is that observed data arises from genuine experience:

$$\forall i,\ y_i \in \Pi_i(\mathcal{A}_{\text{real}}), \tag{3}$$

where $\mathcal{A}_{\text{real}} \subset \mathcal{A}$ denotes the subspace of states realizable by actual world events. Under this assumption, and given sufficient diversity of projections, the system converges to a state that approximately corresponds to real experience.

The entire structure of the analysis to follow concerns what happens when this assumption fails.

Synthetic Projection Contamination

High-fidelity generative systems such as Sora are capable of producing video content that is self-referential, visually coherent, and replayable. When this content depicts the user within familiar scenes, it produces observations $\{y_i^{\text{syn}}\}$ with the following properties. First, each observation is locally admissible: for each i, there exists some $X_i \in \mathcal{A}$ such that $\Pi_i(X_i) = y_i^{\text{syn}}$. The synthetic signal passes each individual projection test. Second, no global preimage exists: there is no $X \in \mathcal{A}_{\text{real}}$ such that $\Pi_i(X) = y_i^{\text{syn}}$ for all i simultaneously. The synthetic observation is not the trace of any real event. Third, the regularizer $\mathcal{R}$ does not strongly penalize these inconsistencies, because the signals are designed to lie within the high-density region of the brain's plausibility manifold.

Definition 1 (Synthetic Projection Contamination)**.** *A reconstruction process is synthetically contaminated if there exists a set of observations $\{y_i^{\text{syn}}\}$ such that:*

1. *(Local admissibility) $\forall i,\ \exists X_i \in \mathcal{A}$ with $\Pi_i(X_i) = y_i^{\text{syn}}$.*
2. *(Global non-realizability) $\nexists X \in \mathcal{A}_{\text{real}}$ with $\Pi_i(X) = y_i^{\text{syn}}$ for all i.*

3. *(Regularizer compatibility) $\mathcal{R}$ does not strongly penalize configurations consistent with $\{y_i^{\text{syn}}\}$.*

The central result of this section establishes that contamination produces stable but false fixed points, and that these are indistinguishable from genuine memories under internal evaluation.

Theorem 1 (Synthetic Projection Contamination Theorem)**.** *Let C be a consistency operator derived from a variational objective with regularizer $\mathcal{R}$. Suppose the observation set contains a non-empty subset of synthetic projections satisfying the contamination conditions above, that $\mathcal{R}$ is insufficiently strict to exclude the contaminated region, and that reconstruction is iterative. Then:*

1. *There exists $X^\dagger \in \mathcal{A} \setminus \mathcal{A}_{\text{real}}$ such that $C(X^\dagger) = X^\dagger$.*
2. *The reconstruction dynamics converge to $X^\dagger$ under repeated application of C with contaminated observations.*
3. *The system cannot internally distinguish $X^\dagger$ from a valid fixed point using projection-level criteria alone.*

Proof sketch. Local admissibility ensures that each constraint $\Pi_i(X) = y_i^{\text{syn}}$ is individually satisfiable. The lack of a global realizable preimage corresponds, in sheaf-theoretic language, to a non-vanishing Čech 1-cocycle over the cover induced by the projection family. However, this cohomological obstruction is invisible to the reconstruction process, which operates locally.

Because $\mathcal{R}$ is not strictly convex over the contaminated feasible region, the energy landscape admits a local minimum $X^\dagger$ that approximately satisfies all projection constraints. Iterative application of C, which performs gradient descent on the variational objective, converges to this local minimum.

Since all evaluation operators Π_i are satisfied to tolerance at $X^\dagger$, no internal diagnostic based on projection residuals can detect the non-realizability. The fixed point $X^\dagger$ is epistemically indistinguishable from a true memory within the system. □

Interpretation: What Sora Does

The article on self-deepfakes describes several phenomena that map cleanly onto this theorem. The "memory twitch"—the growing sense of familiarity with events one knows to be synthetic—corresponds to the system entering the basin of attraction of $X^\dagger$. Repeated viewing corresponds to repeated application of C with contaminated observations, driving convergence. The fading of the "I made this video" tag corresponds to the decay of a constraint that previously excluded the false fixed point.

Most strikingly, the observation that Sora-generated content activates spatial memory structures—that viewers form allocentric cognitive maps of depicted scenes as if they had navigated them—illustrates the depth of the contamination. The synthetic signal is not merely passing a surface plausibility check. It is engaging the geometric encoding layer normally reserved for real-world traversal, activating vector flow fields in the RSVP language corresponding to spatial navigation. This is closer to manifold injection into experiential geometry than to the simple formation of a false belief.

The Key Shift: From Misinformation to Reconstruction Interference

Most discourse about synthetic media is framed around misinformation: false claims about external events propagated through convincing media. This framing misses the deeper problem. The Contamination Theorem describes something more radical: perturbation of the internal reconstruction operator itself.

This is not persuasion. It is not misinformation in the classical sense. It is the injection of synthetic constraints into the closure process that defines the self. In the language introduced earlier: identity is a stable fixed point under recursive constraint closure over autobiographical data. If the constraints are perturbed, the fixed point shifts. The result is not a change in belief but a change in who the agent is.

The Notification Index and the Temporal Cost of Integration

Before proceeding to further theoretical results, we introduce a metric that proves useful across multiple domains. The notification index of a proposed transformation—whether a memory, a Git commit, or a conversational prompt—is the elapsed time Δt between the moment it is proposed and the moment it is accepted, integrated, or definitively rejected.

This quantity decomposes into three components. The first is attention latency Δt_{att}: the time before the proposal enters active consideration. The second is processing latency Δt_{proc}: the time required to render the proposal interpretable. The third is semantic integration latency Δt_{int}: the time required to determine admissibility and incorporate the proposal without producing incompatibilities. Thus

$$\Delta t = \Delta t_{\text{att}} + \Delta t_{\text{proc}} + \Delta t_{\text{int}}. \tag{4}$$

Each component reveals different properties of the receiving system. Attention latency reflects the recency-weighted structure of visibility: in GitHub, proposals not in the recently active part of the repository may simply not be seen; in LLM dialogue, prompts at the beginning of a long conversation exert progressively less influence as recency dominates the context window. Processing latency reflects operational overhead: a data export request, for instance, suspends synchronous interaction entirely, converting a real-time exchange into an asynchronous process governed by event completion. The user grants permission for deferred notification—a form of giving permission to be contacted later—rather than receiving an immediate response. Semantic integration latency reflects the actual difficulty of incorporating the proposed transformation under the system's constraint set.

It is critical to resist reading the notification index as a pure measure of semantic quality. A long delay may reflect not the difficulty of a proposal but simply that the evaluating agent is occupied, that the proposal has fallen out of the recency-weighted attention field, or that operational overhead dominates for reasons independent of content.

The Extended Stack: Synthesis, Entropy, and Rapport

Three additional layers govern whether proposed transformations can be successfully integrated into a shared state. Theoretical synthesis concerns the capacity of a system to reconcile multiple transformations into a single coherent world-state. Two individually admissible changes may be mutually incompatible when composed—corresponding in sheaf-theoretic language to a failure of the gluing condition. Synthesis is the process by which such failures are addressed through reinterpretation, refactoring, or lifting to a higher-order structure.

Entropy budget is the resource constraint on this process. Every proposed transformation carries semantic entropy: ambiguity, degrees of freedom, and unresolved structure that must be resolved before clean incorporation. A minimal semantic change is therefore not only one that isolates a single conceptual effect, but one that does so with minimal entropy. This provides a quantitative justification for why small, precise commits, prompts, and memories are more easily integrated than large, entangled ones: they draw less on the system's capacity for constraint closure.

Rapport is the third layer. It can be understood as a shared compression scheme: a set of aligned priors, conventions, and expectations that allow participants to encode and decode transformations efficiently. In high-rapport settings, a transformation can be specified with relatively little detail because much of its structure is implicitly understood. Rapport acts as a multiplier on the entropy budget, determining how much complexity can be handled without loss of coherence.

These three layers interact. High rapport enables processing of higher-entropy proposals. Strong synthesis reduces effective entropy by reorganizing disparate transformations. Low rapport and weak synthesis force the system to operate under a tighter entropy budget. The notification index, reconsidered in this light, is a composite observable:

$$\Delta t \;\sim\; f(\Delta t_{\text{att}},\; E_{\text{ent}},\; R_{\text{rapport}},\; \mathcal{T}_{\text{synth}})\,, \tag{5}$$

where E_{ent} is the entropy load of the proposal, R_{rapport} the effective compression ratio, and $\mathcal{T}_{\text{synth}}$ the synthesis cost.

Provenance Constraints and Their Non-Robustness

A natural response to synthetic contamination is to mark generated content with provenance signals—watermarks, metadata, cryptographic signatures—that preserve traceability across the content lifecycle. The effectiveness of such mechanisms, however, depends entirely on their invariance under real-world interaction.

Visible watermarks are the weakest form of provenance constraint. They operate at the representational surface and are easily stripped by re-encoding, cropping, or downstream processing. The rapid emergence of third-party watermark removal services demonstrates not a minor loophole but systematic optimization against the constraint: the presence of numerous removal tools indicates that removal is economically incentivized, converting the watermark from a binding constraint to a soft annotation.

Let $\mathcal{T}$ denote the set of transformations induced by routine user interaction and downstream processing. A provenance constraint P is effective only if it satisfies

$$P(X) = P(T(X)) \quad \forall T \in \mathcal{T}. \tag{6}$$

If there exists even a low-cost $T \in \mathcal{T}$ such that $P(T(X)) \neq P(X)$, the constraint fails to define a stable property of the system's outputs. Under such conditions, the system behaves as though the constraint is absent, regardless of its nominal inclusion at the generation stage.

This reframes the significance of the Sora product lifecycle. Two observable features of early self-deepfake interfaces—the absence of synchronized audio and the presence of watermark bypass pathways—reveal not primarily public relations decisions but structural misfits between capability and governance. The absence of audio indicates that the system was not yet operating within a unified multimodal manifold, producing projections that satisfied local constraints (visual plausibility) while failing to enforce cross-modal consistency. The watermark bypassability indicates that provenance enforcement was not embedded in the representation but only annotated onto it.

The transition toward more embedded provenance mechanisms—imperceptible watermarking, cryptographic signing, metadata standards—represents an attempt to produce invariants that persist under a wider class of transformations. But this shift introduces a new tradeoff: as constraints become structurally integrated, they become less legible. The user can no longer directly perceive whether media is synthetic or traceable. Verification is deferred to external systems, shifting the legibility problem upstream rather than solving it.

The general principle is stark:

A constraint is only real if it is invariant under the dynamics of the system it inhabits. (7)

A provenance mechanism that disappears on first contact with the user ecosystem was never part of the operative constraint set. It was decoration on the boundary.

Cryptographic Externalization: Ledgers and Adversarial Robustness

---

The progression from individual memory to institutional record-keeping culminates in systems explicitly designed to withstand adversarial manipulation. A cryptographic ledger is an append-only sequence of states $\{X_t\}$ equipped with a verification function V such that each transition satisfies $V(X_{t+1} \mid X_t) = 1$, and any attempt to alter a past state invalidates all subsequent states. To modify the record is to violate the verification function, rendering the altered state inadmissible.

This represents a maximal extension of externalization, hardened against transformation. Unlike surface-level annotations, cryptographic constraints are embedded in the representation itself. Their robustness is therefore qualitatively different from institutional trust-based records—they derive inviolability from structural dependency rather than social compliance.

However, this robustness relocates rather than eliminates the fundamental problem. A

ledger can faithfully preserve a false or misleading entry, stabilizing an incorrect constraint rather than a correct one. Cryptographic externalization does not guarantee truth; it guarantees persistence of whatever was recorded. The challenge shifts from preserving records to ensuring that what is recorded corresponds to an admissible and meaningful state.

More broadly, this illustrates the hierarchy of externalization: individual vocal encoding augments the constraint set with social observations; written records create persistent artifacts that do not undergo reconsolidation; institutional systems (testimony, contracts, logs) distribute memory across agents and procedures; cryptographic ledgers maximize robustness by enforcing invariance under adversarial conditions.

Each level trades flexibility for constraint strength. Each level also relocates rather than eliminates the problem of grounding.

Collective False Memory: Sheaf Obstruction Across Agents

The dynamics of reconstruction extend across groups, where they can produce coordinated yet non-veridical outcomes. Documented reports surrounding the apparitions at Our Lady of Medjugorje include instances in which large numbers of individuals converge on descriptions of events that are not independently corroborated. These can be understood as emergent properties of coupled reconstruction systems rather than theological or sociological anomalies alone.

At smaller scale, the same mechanism is observable in ordinary social settings. When a salient description is introduced—particularly by a trusted or emotionally authoritative speaker—it functions as a constraint that propagates rapidly across agents. Individuals who did not directly perceive the described event may nonetheless incorporate it into their reconstruction, especially if it aligns with ambient expectations. A child's verbal report of "red eyes outside the window" can trigger near-immediate agreement among peers who were not attending to the window, and even among those whose visual processing was otherwise engaged. The verbalized description acts as a high-salience constraint, and in the absence of strong contradictory input, it is incorporated into multiple recon-

structions simultaneously. What emerges is a shared memory that presents as jointly observed, despite originating from a single verbal source.

Formally, let agents $\{A_k\}$ have reconstruction spaces with respective projection families $\{\Pi_i^{(k)}\}$. The introduction of a shared signal y_s augments each agent's constraint set:

$$\{\Pi_i^{(k)}\} \;\rightarrow\; \{\Pi_i^{(k)}\} \cup \{\Pi_s^{(k)}(X) = y_s\}. \tag{8}$$

If y_s is sufficiently plausible and not contradicted by strong sensory evidence, each agent may converge toward a reconstruction that includes it, yielding synchronized false memories without any underlying event.

The multi-agent failure is more precisely expressed in sheaf-theoretic terms. Let agents A and B have reconstruction spaces with a shared overlap domain U_{AB}. If A incorporates synthetic observations and B does not, then A reconstructs $X_A^\dagger \notin \mathcal{A}_{\text{real}}$ while B reconstructs $X_B^* \in \mathcal{A}_{\text{real}}$. On the overlap:

$$\rho_{AB}(X_A^\dagger) \neq \rho_{BA}(X_B^*), \tag{9}$$

defining a non-trivial cohomology class $[\alpha] \in H^1(U, \mathcal{S}) \neq 0$. No global shared memory exists. The "remember when?" failure described in accounts of collective false memory is literally a failure of global section existence over a shared domain.

Remark. *Collective agreement is not evidence of truth. It is evidence of constraint alignment. Two systems can synchronize on a false state just as readily as on a true one, if the propagating constraint is sufficiently plausible and the competing constraints sufficiently weak.*

Philip K. Dick and the Inversion of Record and Experience

Philip K. Dick's persistent literary concern—the instability of reality when memory and record diverge—anticipates the formal problem with unusual clarity. In *Minority Report*,

the concept of precognition is often read as a speculative ability to see the future. More precisely, the depicted system produces authoritative records that precede experience, inverting the usual temporal relationship between observation and documentation. The record ceases to be a trace of reality and becomes a generator of it.

This inversion is now technically realizable. Synthetic media systems produce documents of events before—or instead of—any actual experience. What Dick imagined as a science-fictional extreme is now a consumer product feature.

Dick's own extensive *Exegesis* was catalyzed partly by a series of anomalous experiences following medication administered during a dental procedure, and has been discussed alongside other altered states—such as those induced by LSD—in which perception, memory, and interpretation become unusually labile. From the reconstruction-theoretic perspective, such episodes can be understood as instances in which the balance of the consistency operator is perturbed: altered neurochemical states may modify the weighting of constraints, reduce the influence of prior regularization, or increase the salience of internally generated signals relative to externally grounded ones. The boundary between admissible and inadmissible configurations is not fixed but depends on the stability of the underlying constraint architecture.

This does not require evaluating the content of such experiences. The relevant point is structural: changes in constraint weighting—whether from neurochemical alteration or from synthetic media injection—can produce the same class of reconstruction failure. The mechanism is the same. The substrate differs.

Externalized Memory and the Role of External Records

Human recollection is not retrieval from a stable archive but reconstruction from partial cues. This reconstruction is inherently non-veridical: there is no privileged "original" to which recall can be compared. Each act of remembering produces a new instantiation of the past, subject to drift, interference, and reinterpretation.

Against this background, external records acquire structural significance beyond their

mnemonic convenience. Physical artifacts—drawings, notes, photographs, ephemera—serve as fixed constraints that do not undergo reconsolidation. They anchor reconstruction by providing invariant reference points that can be revisited and compared across time. Crucially, these artifacts are not only aids to individual recall but instruments of social validation. Shared reality depends on establishing agreement across agents, and such agreement is mediated through externalized evidence rather than internal confidence.

An individual who recognizes early that internal memory is insufficient for establishing shared reality—who understands that the felt conviction of a recollection carries no epistemic authority in the absence of external corroboration—will develop systematic practices of externalization. This is not a compensation for deficient memory but a recognition that only the externalized regime reliably participates in the construction of shared facts.

There is a distinction between two memory regimes: one relying on internal reconstruction, efficient but susceptible to contamination; the other anchored in external artifacts, slower and less flexible but constraint-stabilized. What Sora and related systems threaten is the integrity of the second regime, by producing artifacts that have the surface properties of genuine external records—replayable, self-referential, detailed—while lacking any correspondence to an actual event.

Institutional Externalization: Testimony, Contracts, and Logs

Legal testimony is not merely the recollection of an event but its formal inscription into a collectively recognized record. The act of speaking under oath introduces social and procedural constraints: the statement is witnessed, transcribed, and subject to cross-examination. These layers transform a private reconstruction into a publicly contestable artifact.

Contracts are forward-looking memory structures that encode agreements in a form persisting independently of the participants' subsequent recollections. By fixing terms in written language, contracts reduce the degrees of freedom available for reinterpretation, constraining future reconstructions of what was agreed. System logs are the most granular instantiation: append-only traces of state transitions with timestamps and metadata, designed for audit and replay.

The common structure across these institutions is this: externalization increases the dimensionality and persistence of the constraint set, reducing the space of admissible reconstructions and promoting convergence toward more stable fixed points. The stability of shared reality depends on the integrity of these artifacts and the processes that maintain them.

Semantic Graph Encoding and the Role of Denouement

The preceding analysis implies that memory stability is a function of constraint richness. Different encoding strategies impose different constraint structures, with direct consequences for long-term fidelity.

One encoding regime privileges imagistic and narrative simulation, in which memories are reconstructed through sensory reactivation and internal storytelling. This regime is efficient but susceptible to drift: each act of recall reintroduces variability through associative blending and contextual reinterpretation. A different regime emphasizes semantic graph encoding, in which experiences are organized as networks of relations, constraints, and logical dependencies. Let $G = (V, E)$ where vertices V represent entities, events, or propositions and edges E encode causality, temporal order, and implication. Recall corresponds to reconstructing a consistent traversal or subgraph rather than reactivating a sensory trace.

The role of denouement in this framework is structural, not aesthetic. A denouement is the point at which causal and logical dependencies are resolved into a closed configuration. When an event is encoded with a clear denouement, its structure becomes constrained by necessity: the relationships between components are fixed by the requirement that the sequence resolves coherently. This closure reduces the degrees of freedom available during recall.

From the perspective of reconstruction theory, semantic graph encoding with denouement strengthens the regularization term $\mathcal{R}$ by embedding tighter relational constraints within the representation itself. The memory is not preserved as a static record but as a constrained solution that remains invariant under admissible transformations. This does

not render memory immutable—errors can still occur through omission or misclassification—but once a semantic graph with a well-defined denouement is established, subsequent recall tends to preserve its structural relations even as peripheral detail degrades.

Formal Encoding and the Minimization of Interpretive Degrees of Freedom

---

If a piece of information is to persist with high fidelity across time and agents, it must be encoded in a form that minimizes the degrees of freedom available during reconstruction. Informal descriptions admit multiple plausible interpretations and are therefore susceptible to drift. Formal representations impose strict constraints on admissible interpretations.

A mathematical formula functions as a highly compressed constraint system. Its symbols are not merely descriptive but operational: they define relations that must hold under all valid interpretations. Small deviations tend to produce contradictions rather than alternative plausible meanings, making corruption more detectable. From the reconstruction-theoretic perspective, formal encoding strengthens $\mathcal{R}$ by embedding invariants directly into the representation, shifting memory from a regime of approximate reconstruction to one of constrained derivation.

The claim is not that only mathematical content can be preserved reliably, but that stability increases with the degree to which encoding constrains interpretation. Legal language, formal specifications, and structured records all approximate this principle to varying degrees. Mathematics represents a limiting case in which ambiguity is minimized and invariance is maximized.

Thus, if the objective is to ensure that a piece of information remains recoverable with minimal distortion, it should be externalized in a form that tightly constrains its reconstruction. The closer an encoding is to a formal constraint system, the more resistant it is to drift across time, agents, and contexts.

The Reconstruction Stability Theorem

We are now in a position to state the central theorem that integrates the preceding analysis.

Theorem 2 (Reconstruction Stability Theorem)**.** *A reconstructed state X^* is stable under temporal evolution and inter-agent communication if and only if the associated constraint system satisfies the following four conditions:*

1. ***Constraint Sufficiency.*** *The combined projection set $\{\Pi_i\}$ is rich enough to reduce the feasible set to a bounded region in $\mathcal{A}$.*

2. ***Constraint Robustness.*** *The constraints are invariant under admissible transformations $T \in \mathcal{T}$:*
$$\Pi_i(X) = y_i \;\Rightarrow\; \Pi_i(T(X)) = y_i.$$

3. ***Constraint Externalization.*** *A subset of constraints is encoded in persistent external artifacts that do not undergo reconsolidation.*

4. ***Constraint Closure.*** *A global solution satisfying all constraints simultaneously exists:*
$$\bigcap_i \Pi_i^{-1}(y_i) \neq \varnothing, \quad X^* \in \mathcal{A}_{\text{real}}.$$

Proof sketch. If all four conditions hold, the feasible set of reconstructions is both restricted and stable under iteration. Sufficiency bounds the solution space; robustness ensures invariance under system dynamics; externalization introduces non-decaying constraints; closure guarantees the existence of a globally consistent, realizable solution. Under these conditions, C admits a unique or stable fixed point X^*, and reconstruction converges.

Conversely, if any condition fails, instability arises. Insufficient constraints yield underdetermination; lack of robustness allows drift under transformation; absence of externalization permits unconstrained reconsolidation; failure of closure produces incompatible local solutions or non-realizable fixed points. □

The theorem subsumes the phenomena analyzed throughout this work as specific failure or success modes.

Corollary 3. *Each of the following is an instantiation of the theorem:*

1. ***Internal Reconstruction Drift.*** *Robustness fails under recall dynamics when memory relies primarily on internal projections without external constraints.*
2. ***Semantic Graph Encoding.*** *Structured relational encoding with denouement closure strengthens $\mathcal{R}$ and reduces admissible reconstructions.*
3. ***Externalization via Record.*** *Speech, writing, and logging augment the projection set with persistent constraints, probabilistically stabilizing reconstruction.*
4. ***Institutional Memory.*** *Testimony, contracts, and logs implement large-scale externalization through redundancy and procedural enforcement.*
5. ***Cryptographic Ledgers.*** *These maximize robustness by enforcing invariance under adversarial transformation, but remain dependent on correct initial encoding.*
6. ***Collective False Memory.*** *Shared constraints introduced across agents without corresponding realizable states satisfy local consistency but violate closure.*
7. ***Synthetic Media Contamination.*** *High-fidelity synthetic observations satisfy projection constraints while lacking a global preimage, breaking injectivity and producing false fixed points.*
8. ***Formal Encoding.*** *Mathematical representations minimize interpretive degrees of freedom, approximating maximal constraint closure within symbolic systems.*

Residual Energy, Attention Gating, and the Zeigarnik Effect

We now introduce a functional that gives the abstract notion of “constraint incompleteness” a precise quantitative form, and connects it to observed phenomena of memory persistence and attention.

Definition 2 (Residual Energy)**.** *Let C be the consistency operator and X_t the system state at time t. Define the residual energy*

$$E_{\text{res}}(X_t) = \|C(X_t) - X_t\|^2, \tag{10}$$

measuring the deviation from a fixed point. When $E_{\text{res}}(X_t) = 0$, the system has achieved closure; when $E_{\text{res}}(X_t) > 0$, there remain unsatisfied constraints.

The relationship between residual energy and memory persistence can be expressed by postulating that recall probability is an increasing function of residual energy over an intermediate regime:

$$P_{\text{recall}}(X_t) \;\propto\; f\big(E_{\text{res}}(X_t)\big), \tag{11}$$

where f is monotonically increasing for moderate values but may saturate or decay if the residual becomes too large, corresponding to incoherent or poorly integrated structures. The accumulated residual over an interval $[t_0, t_1]$ is

$$\mathcal{E}(X) = \int_{t_0}^{t_1} E_{\text{res}}(X_t)\, dt, \tag{12}$$

providing a measure of total "unfinishedness" over that period.

Attention can be modeled as a weighting mechanism driven by this residual energy. Let $w_i(t)$ denote the attentional allocation to constraint i. Then

$$w_i(t+1) = w_i(t) + \alpha \frac{\partial E_{\text{res}}}{\partial \Pi_i}, \tag{13}$$

so that constraints contributing most to the residual receive increased focus. This aligns with the observation that unresolved elements attract disproportionate cognitive resources—the phenomenon standardly described as the Zeigarnik effect. Incomplete tasks persist in memory because they sustain nonzero residual energy, driving continued iteration of the reconstruction process.

Closure corresponds to the dissipation of residual energy. When a denouement is intro-

duced, the system transitions to a fixed point X^* with $E_{\text{res}}(X^*) = 0$. Attentional weights decay, but the representation remains stable due to the constraints accumulated during the period of elevated residual energy.

Nested Scope, Interrupted Narrative, and Encoding Amplification

Igor Ledochowski's treatment of nested interrupted storylines in conversational hypnosis provides a structurally interesting application of the residual energy framework. When a narrative introduces an unresolved structure—an open scope—the system maintains a nonzero residual corresponding to the outstanding constraint. Subsequent content is processed in the presence of this residual, increasing its salience and integration. When the original narrative is eventually resumed and resolved, the system undergoes closure, satisfying the outstanding constraint.

In formal terms: let a narrative segment introduce a partial structure with unresolved dependencies. These dependencies constitute open constraints that sustain elevated E_{res} across multiple processing steps. When closure is eventually achieved, the transition from elevated residual to zero residual consolidates the entire accumulated constraint set into a single stable fixed point. The embedded content is thereby encoded more deeply and more durably than it would be under immediate closure.

This mechanism explains why interrupted endings, nested storylines, and "almost finished" states are disproportionately memorable. They do not merely leave a task incomplete; they create a sustained period of constraint maintenance that forces deep structural integration prior to resolution.

Lemma 4 (Interrupted Closure Amplification)**.** *Let $\{X_t\}$ be a reconstruction trajectory under C, and let $E_{\text{res}}(X_t)$ be the residual energy. Suppose the system receives a closure signal at time t_0—an indication that a fixed point is imminent—but closure is not actually achieved. Then for $t > t_0$, the system exhibits a transient increase in residual energy and attentional allocation, resulting in enhanced persistence and continued iteration of the reconstruction process.*

Proof sketch. A closure signal induces a shift in system dynamics toward convergence, lowering expected residual energy $\mathbb{E}[E_{\text{res}}(X_{t_0+1})]$. If closure is not realized, the dis-

crepancy between expected and actual state introduces a secondary residual $\Delta E = E_{\text{res}}(X_{t_0}) - \mathbb{E}[E_{\text{res}}(X_{t_0})]$. This mismatch amplifies attentional weights: $w_i(t+1) = w_i(t) + \alpha \Delta E$. The trajectory remains in a non-fixed-point regime, maintaining or strengthening the representation rather than allowing it to decay. □

Corollary 5 (Conversational Persistence)**.** *In interactive systems, the introduction of an apparent termination signal followed by continuation increases engagement without requiring additional informational content. The effect arises purely from the manipulation of constraint timing.*

Remark. *The observation that concluding a conversation with "ok thanks that's good enough for now" before continuing—rather than simply continuing—produces a detectable effect on subsequent engagement is a direct application of this lemma. By suggesting closure and then not delivering it, the system creates a transient residual spike. No new information is introduced; only the temporal profile of constraint satisfaction is modified. The same principle underlies serial fiction, thesis chapters that end on questions, and academic papers whose introductions promise theorems not yet delivered.*

Closure Without Grounding

The Contamination Theorem established that stable false memories exist. The following proposition generalizes this result to a cross-substrate principle.

Proposition 6 (Closure Without Grounding)**.** *Let $\mathcal{A}_{\text{real}}$ denote the subset of admissible states corresponding to grounded, externally realizable histories. A system exhibits closure without grounding if there exists $X^* \in \mathcal{A}$ such that $C(X^*) = X^*$ but $X^* \notin \mathcal{A}_{\text{real}}$. Such states are internally stable yet lack a valid generative origin.*

This condition arises whenever the constraint system is sufficient to enforce internal consistency but insufficient to guarantee correspondence with external reality. It appears in structurally identical form across three substrates.

In dream states, external sensory projections are suppressed and reconstruction operates

over internally generated constraints. The system converges to a fixed point satisfying autobiographical and semantic coherence, producing a state experienced as memory. The conviction that a painting encountered in a dream had been previously created, or that one possesses an impossible anatomy—an additional thumb revealed beneath a fingernail—reflects the achievement of internal closure without grounding. The body schema has satisfied its internal consistency conditions while violating the structural invariants of anatomical topology. Upon reintroduction of external constraints on waking, the non-realizability becomes apparent.

In social systems, a shared constraint propagates across agents, leading to synchronized reconstruction of a state locally consistent for each agent but globally unrealizable. The system achieves inter-agent closure without correspondence to an underlying event.

In large language models, outputs are produced by satisfying learned statistical and semantic constraints over a latent manifold. When the constraint system lacks sufficient grounding signals, the model generates outputs that are internally coherent but do not correspond to real-world facts. These outputs are fixed points in the model's constraint space that fail to map to grounded states. CLIO faces a related failure mode: it assumes that inconsistencies signal missing structure to be reconstructed. Under contamination, however, the inconsistency is irreducible. The system responds by over-smoothing, collapsing resolution and admitting false fixed points. The misclassification of synthetic obstruction as reconstructible incompleteness is precisely the CLIO failure mode identified by the Contamination Theorem.

Corollary 7 (Indistinguishability Under Internal Evaluation)**.** *No procedure operating solely on internal constraints can reliably distinguish grounded from ungrounded states. External validation is necessary to resolve this ambiguity.*

Encoding Modality and Constraint Precision

Different input modalities impose systematically different constraint structures on the encoding process, leading to predictable differences in the fidelity of reconstructed content.

High-precision modalities such as keyboard typing—whether QWERTY, Dvorak, or gesture-based systems like Swype—enforce strong orthographic and sequential constraints. Each character or gesture is explicitly selected, and errors are local, detectable, and correctable. The resulting projection closely preserves the intended representation, yielding a narrow feasible set.

Vocal encoding followed by transcription introduces a noisier procedure $\tilde{P}$, in which phonetic approximation, segmentation ambiguity, and model-based inference shape the output. The corresponding projection is perturbed: $\tilde{\Pi} = \Pi + \epsilon$, where ϵ represents modality-induced noise. Reconstruction under these conditions produces outputs that are often semantically plausible but not necessarily faithful to the intended structure.

This explains the observed substitution of "ecological recall" for "ecphoric recall" under transcription conditions. The former is not a random error but a locally optimal reconstruction under degraded constraints: phonetically similar, semantically coherent, lexically familiar. It satisfies all projection-level admissibility criteria for the noisy procedure while diverging from the intended concept.

Remark. *A system does not reproduce what was intended; it produces what is admissible under the constraints of its encoding procedure. Errors arising from encoding modality shifts are not random but structured: they occupy the nearest admissible point under the degraded constraint set.*

The correction of such an error—by explicitly reintroducing the intended term—constitutes a reapplication of stronger external constraints, collapsing the residual ambiguity:

$$\tilde{\Pi} \;\to\; \Pi \;\Rightarrow\; X'^{*} \to X^{*}. \tag{14}$$

Typing introduced micro-delays and discrete selection steps that allow partial constraint evaluation during construction, embedding a form of incremental closure. Speech-to-text encoding is more continuous and less interruptible, increasing reliance on post hoc correction. The choice of input modality is therefore not merely ergonomic but epistemic: it determines the structure of the constraint system through which thought is externalized.

Provenance Ambiguity and Ecphoric Source Reconstruction

A recurring feature of everyday cognition is difficulty recalling not only what one knows but how that knowledge was acquired. This arises because content and provenance are governed by partially distinct constraint sets. Semantic coherence, internal consistency, and integration with prior knowledge constrain the content reconstruction X^*, whereas the encoding pathway—whether information was read fully, skimmed, heard second-hand, or inferred—leaves only weak or decaying traces.

Formally, let S denote the source of an encoded state. While the system may converge on a stable X^*, the corresponding source reconstruction remains underdetermined:

$$C_{\text{content}}(X^*) = X^*, \quad C_{\text{source}}(S) \not\Rightarrow S^* \text{ unique}. \tag{15}$$

In ecological recall conditions—where content is encountered in fragmented, repeated, or transformed forms across multiple channels—each encounter contributes constraints to X^* while diluting constraints associated with S. The term "ecological recall" is therefore not entirely incorrect as a description of real-world memory behavior, even if it is not the technical term intended. The inadvertent coinage captures something real about the multi-source, diffuse nature of how knowledge accumulates.

Remark. *Knowledge can be stable while its origin is not. Reconstruction preserves structure more readily than it preserves lineage.*

The Synchronization Principle

The final result of this essay concerns the dynamics of constraint systems across multiple instances that share a common origin but are subsequently decoupled.

The discontinuation of Swype—the gesture-based keyboard system—provides a concrete illustration. Initially, multiple devices share a common mapping P_0 from gesture trajectories to lexical outputs, defining a shared constraint system $\mathcal{C}_0$ that ensures consistent

reconstructions across contexts. Once centralized updates cease, each device accumulates local adaptations based on usage patterns and corrections:

$$P_0 \;\to\; P_t^{(1)},\; P_t^{(2)},\; \ldots, \quad \mathcal{C}_0 \;\to\; \mathcal{C}_t^{(1)},\; \mathcal{C}_t^{(2)},\; \ldots \tag{16}$$

The same gesture trajectory may yield different outputs depending on the device. The system remains internally consistent on each device but loses global consistency across devices. Personalization, while locally beneficial, produces fragmentation at the global level.

Theorem 8 (Synchronization Principle)**.** *Global alignment across instances $\{(k)\}$ initialized from a common procedure P_0 is preserved if and only if there exists a synchronization operator S such that*

$$S\Big(P_t^{(1)}, \ldots, P_t^{(n)}\Big) \;\approx\; P_{t,\text{shared}}, \tag{17}$$

applied with sufficient frequency relative to the rate of local adaptation.

Proof sketch. Local updates introduce idiosyncratic perturbations $\delta^{(k)}$. In the absence of synchronization, these accumulate independently, and $\|P_t^{(i)} - P_t^{(j)}\| \to \infty$ as $t \to \infty$. A synchronization operator reprojects individual procedures onto a shared manifold, bounding divergence: $\|P_t^{(i)} - P_t^{(j)}\| \leq \epsilon$. Thus alignment depends not on initial conditions but on ongoing coordination. □

This principle instantiates across scales: in human memory and communication, shared narratives and repeated interaction function as synchronization mechanisms; in machine learning systems, independently fine-tuned models diverge unless periodically synchronized; in cultural systems, institutions such as education, media, and publication act as synchronization operators maintaining mutual intelligibility across populations.

Remark. *Systems do not remain aligned by sharing an origin; they remain aligned by sharing an ongoing process of synchronization. Drift is the default. Alignment is an active process.*

Latent Externalization and Version-Controlled Constraint Layers

Version control systems introduce a distinct form of externalization in which information is rendered inactive under default reconstruction procedures rather than deleted. When sections of a document are removed from the visible artifact but preserved in revision history, they remain within the broader constraint system while ceasing to influence standard interpretation.

Let $\mathcal{C}_{\text{full}}$ denote the total constraint set encoded across all versions and $\mathcal{C}_{\text{vis}} \subset \mathcal{C}_{\text{full}}$ the subset present in the current visible state. Reconstruction by a typical reader proceeds under $\mathcal{C}_{\text{vis}}$, while reconstruction by an agent traversing version history operates under an expanded set approaching $\mathcal{C}_{\text{full}}$. The hidden sections constitute latent constraints: elements of the system that are not active in the default reconstruction pathway but remain recoverable under alternative procedures.

Unlike erased information, which is removed from the constraint system entirely, latent constraints persist and can be reactivated without loss of fidelity. This creates a layered constraint architecture: the visible document defines the primary constraint surface; version history defines a secondary layer encoding alternative paths, discarded hypotheses, and intermediate structures.

Remark. *Information that is preserved but not accessible under the default reconstruction procedure functions as a latent constraint: it exists within the system but does not participate in the formation of the current fixed point. What exists in a system is not the same as what constrains it.*

Procedural Control and the Induction of Constraint Systems

Reconstruction can be influenced not only by modifying observations or constraints but by shaping the procedures through which constraints are generated and evaluated. This is a higher-order mechanism: instead of altering y_i or Π_i directly, one alters the procedure P that produces $\mathcal{C}$.

Let $P : \mathcal{I} \to \mathcal{C}$ denote a procedure mapping raw interaction into an effective constraint set. Under this formulation, reconstruction becomes a two-stage process:

$$\mathcal{I} \xrightarrow{P} \mathcal{C} \xrightarrow{\text{closure}} X^*. \tag{18}$$

A change in procedure $P \to P'$ induces a corresponding change in constraint system $\mathcal{C} \to \mathcal{C}'$, which may admit different fixed points even when the underlying inputs $\mathcal{I}$ remain unchanged.

The conviction that there are eight days in a week—arising under LSD from a counting procedure that includes both endpoints—illustrates this mechanism with unusual clarity. The system does not violate constraints to reach this conclusion; it operates under a modified procedure P' that generates a constraint set admitting a different cardinality. The accompanying conviction is real because the system evaluates truth relative to the active constraint set. Under P', the standard seven-day structure appears artificially constrained. Cultural artifacts—song titles, calendar patterns—are reinterpreted as corroborating the modified framework.

The suggestion that this conviction could be induced in others by guiding them through a particular counting procedure illustrates constraint induction via procedural framing. By structuring the process through which elements are enumerated, one can temporarily shift the effective constraint system of another agent, leading them to converge on the same locally consistent but globally incorrect result.

Proposition 9 (Procedural Induction of Fixed Points)**.** *Given a reconstruction system with consistency operator C, there exist procedures P and P' such that, for the same input $\mathcal{I}$, the induced constraint sets $\mathcal{C}$ and $\mathcal{C}'$ yield distinct fixed points $X^* \neq X'^*$. Moreover, if P' is constructed such that $\mathcal{C}'$ remains internally consistent, then X'^* will be experienced as equally valid within the system.*

The most effective way to change what a system concludes is not to change the conclusion, nor even the constraints, but to change the procedure that generates the constraints. This insight is central to many practices—narrative framing, guided questioning, prompting strategies for language models—in which the form of engagement shapes the content of the result.

Conclusion: Accumulation Requires Closure, Closure Requires Grounding

---

The central thesis of this essay can be stated in three levels.

At the phenomenological level: memory is not storage. It is iterative reconstruction under partial constraints. Its stability is a function of constraint richness, not storage fidelity.

At the formal level: the Reconstruction Stability Theorem establishes that stability requires sufficiency, robustness, externalization, and closure. The Contamination Theorem establishes that high-fidelity synthetic observations can produce stable but non-realizable fixed points that are indistinguishable from genuine memories under internal evaluation. The Closure Without Grounding proposition unifies this failure mode across dream states, collective false memory, and LLM hallucination. The Synchronization Principle establishes that alignment across instances decays without active coordination.

At the structural level: truth in a reconstruction-based system is not correspondence to a stored record. It is survival as a fixed point under a sufficiently rich, robust, and externalized constraint closure process—with the additional requirement that the fixed point be realizable, i.e., that it lie in $\mathcal{A}_{\text{real}}$.

The distinction between closure and grounding is the deepest result. A system can achieve perfect internal consistency—every constraint satisfied, every projection matched, every narrative coherent—and still be wrong. What distinguishes veridical reconstruction is not the absence of residual error but the presence of constraints that tie the reconstruction to an external generative process. Systems that optimize only for closure will inevitably produce coherent but non-realizable outputs. Systems that incorporate grounding constraints can align reconstruction with reality, but only to the extent that those constraints are robust, persistent, and correctly specified.

As generative systems approach the fidelity required to produce admissible autobiographical projections, and as those projections are integrated into unified multimodal

interfaces that blur the boundary between tool use and experience, the primary challenge shifts: from improving realism to enforcing constraint grounding. Without mechanisms that guarantee cross-modal coherence, persistent provenance, and alignment with realizable states, synthetic signals propagate through cognitive and social processes as if they were unmarked components of reality.

The residual energy functional offers a final observation. What persists in memory is not what is observed but what remains unresolved long enough to be structurally integrated and then closed. Systems that exploit this—by introducing constraints that sustain elevated residual energy prior to resolution—shape not just what is remembered but what is deeply encoded. That this principle applies equally to narrative craft, conversational strategy, and synthetic media injection should give pause. The same mechanism that supports effective teaching, meaningful storytelling, and coordinated cognition is also the mechanism by which reconstruction can be most deeply contaminated.

Stability is not truth. Closure is not grounding. Accumulation without constraint closure is noise. Constraint closure without grounding is hallucination.

Note on CLIO and TARTAN. The failure mode described in the Contamination Theorem maps directly onto a known vulnerability in the CLIO reconstruction framework. CLIO assumes that detected inconsistencies signal missing structure to be recovered through further constraint application. Under synthetic contamination, however, the inconsistency is irreducible: no additional reconstruction can produce a globally realizable state, because no such state exists. The system's response—over-smoothing, collapsing resolution, accepting the false fixed point—is exactly the CLIO failure mode identified by the theorem. A TARTAN-level response would require a strengthened regularizer $\mathcal{R}^+$ capable of distinguishing reconstructible incompleteness from synthetic obstruction, paired with a detection criterion for constraint sets that are locally admissible but globally non-realizable. This remains an open problem, and arguably the most important open problem in the formal theory of reconstruction under adversarial conditions.

References

[1] *Elizabeth F. Loftus. Eyewitness Testimony. Harvard University Press, 1979.*

[2] *Elizabeth F. Loftus and John C. Palmer. Reconstruction of Automobile Destruction: An Example of the Interaction Between Language and Memory. Journal of Verbal Learning and Verbal Behavior, 13(5), 1974.*

[3] *Lynn Nadel and Morris Moscovitch. Memory Consolidation, Retrograde Amnesia and the Hippocampal Complex. Current Opinion in Neurobiology, 7(2), 1997.*

[4] *Karim Nader, Glenn E. Schafe, and Joseph E. LeDoux. Fear Memories Require Protein Synthesis in the Amygdala for Reconsolidation After Retrieval. Nature, 406, 2000.*

[5] *Frederic C. Bartlett. Remembering: A Study in Experimental and Social Psychology. Cambridge University Press, 1932.*

[6] *Daniel L. Schacter. The Seven Sins of Memory: How the Mind Forgets and Remembers. Houghton Mifflin, 2001.*

[7] *Henry L. Roediger III and Kathleen B. McDermott. Creating False Memories: Remembering Words Not Presented in Lists. Journal of Experimental Psychology: Learning, Memory, and Cognition, 21(4), 1995.*

[8] *Bluma Zeigarnik. Über das Behalten von erledigten und unerledigten Handlungen. Psychologische Forschung, 9, 1927.*

[9] *Endel Tulving. Elements of Episodic Memory. Oxford University Press, 1983.*

[10] *Martin A. Conway. Episodic Memory and the Self. In Memory, Brain, and Belief, Harvard University Press, 2000.*

[11] *Karl Friston. The Free-Energy Principle: A Unified Brain Theory? Nature Reviews Neuroscience, 2010.*

[12] *Andy Clark. Surfing Uncertainty: Prediction, Action, and the Embodied Mind. Ox-*

ford University Press, 2016.

[13] Masaki Kashiwara and Pierre Schapira. Sheaves on Manifolds. Springer, 1990.

[14] David I. Spivak. Category Theory for the Sciences. MIT Press, 2014.

[15] Saunders Mac Lane. Categories for the Working Mathematician. Springer, 1971.

[16] Roger Penrose. The Emperor's New Mind. Oxford University Press, 1989.

[17] Philip K. Dick. The Exegesis of Philip K. Dick. Houghton Mifflin Harcourt, 2011.

[18] Philip K. Dick. The Minority Report. 1956.

[19] Ashish Vaswani et al. Attention Is All You Need. NeurIPS, 2017.

[20] Tom B. Brown et al. Language Models are Few-Shot Learners. NeurIPS, 2020.

[21] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. Deep Learning. MIT Press, 2016.

[22] Andrej Karpathy. The Unreasonable Effectiveness of Recurrent Neural Networks. Blog post, 2015.

[23] Igor Ledochowski. Conversational Hypnosis. Hypnosis Training Academy, 2009.

[24] Claude E. Shannon. A Mathematical Theory of Communication. Bell System Technical Journal, 1948.

[25] Thomas M. Cover and Joy A. Thomas. Elements of Information Theory. Wiley, 2006.

[26] James Gurney. Dinotopia: A Land Apart from Time. Turner Publishing, 1992.

[27] Alexander G. Huth et al. Natural Speech Reveals the Semantic Maps That Tile Human Cerebral Cortex. Nature, 2016.

[28] Elena Piech. Research on Spatial Memory Encoding. Various publications.

[29] Satoshi Nakamoto. Bitcoin: A Peer-to-Peer Electronic Cash System. 2008.

[30] Whitfield Diffie and Martin E. Hellman. New Directions in Cryptography. IEEE

Transactions on Information Theory, 1976.

[31] Luciano Floridi. The Philosophy of Information. Oxford University Press, 2011.