Compressed Causality

Stability, Projection, and the Geometry of Alternative Histories

Flyxion

2026

For every system that must act in the present
on the basis of a history too large to store.

Preface

Every system that survives must forget.

A mind cannot retain every experience. An institution cannot preserve every event. A scientific theory cannot represent every detail of every observation. Even the observable universe contains only a fragment of the information contained in its own history.

This fact is so familiar that it is easy to overlook its significance. Every act of remembering presupposes forgetting. Every representation presupposes compression. Every record of the past is already a selection from a history too large to preserve in full.

The central question of this book emerges from that observation.

If the past is compressed, when can it still be recovered?

More precisely: under what conditions can a system reconstruct answers to questions about alternative histories from a representation that does not contain those histories explicitly?

This question appears in many forms. Machine learning systems must answer questions about training data that no longer exists at inference time. Scientific theories must account for observations that vastly exceed what any single model can represent. Legal institutions must reason about social realities from archives that preserve only fragments of them. Cosmologists must infer the history of the universe from signals that survive in the present configuration of matter, radiation, and geometry.

Although these domains appear unrelated, they share a common structure. Each begins with a space of possible histories—a space of trajectories through which the system might have passed. Each applies a compression that preserves some distinctions and discards others. Each attempts to answer questions about trajectories that are no longer directly accessible.

The framework developed in this book studies that shared structure.

Part One develops the mathematics of reconstruction from compressed representations. Part Two examines three domains—learning systems, semantic inference, and physical field dynamics—in which reconstruction can be proved or rigorously tested under explicit conditions. Part Three extracts a diagnostic methodology applicable to any domain where reconstruction is claimed or required. Part Four applies that methodology to memory, institutions, and cosmology. Part Five asks what these results imply about explanation, representation, and the geometry of alternative histories.

The goal throughout is not to prove that reconstruction is always possible. In many systems it is not, and the framework is designed to identify precisely when and why it fails. Nor is the goal to defend any particular theory of memory, cognition, cosmological dynamics, or institutional design. The goal is narrower and, I believe, more fundamental: to identify the conditions under which reconstruction from compressed representations becomes possible at all, and to characterize precisely what happens when those conditions are not met.

Those conditions turn out to be three. The history's influence on the query must decay with interaction order—the stability condition. The compression must preserve the information the query requires—the sufficiency condition. And the compression must be calibrated to the stability structure so that both conditions are controlled by a single depth parameter—the alignment condition. When all three hold, the present serves as a recoverable record of the past. When any one fails, it fails in a specific, diagnosable way that the framework can characterize precisely.

The chapters that follow develop this claim from its mathematical foundations to its philosophical consequences.

They begin with compression.

They end with memory.

Between those two ideas lies a single question: what kind of world allows its history to remain recoverable after it has been compressed?

This book is an attempt to answer that question.

Reader's Guide

This book develops a single framework across multiple levels of abstraction, from formal mathematical results to applications in cognition, institutions, and cosmology. Readers interested in different aspects of the project may wish to approach the material through different paths.

The Structure of the Book

The book is organized into five parts.

Part One develops the mathematical foundation. It introduces trajectory spaces, compression operators, interaction decompositions, stability sequences, reconstruction operators, and the Fundamental Error Decomposition. Readers interested primarily in the formal structure of the framework should begin here. The four chapters of Part One are self-contained and do not presuppose familiarity with any of the application domains.

Part Two instantiates the framework in three domains where the conditions for reconstruction can be established rigorously: learning systems, semantic inference, and physical field systems. These chapters develop the framework's verified classes, apply the five diagnostic questions to each domain, and identify the open empirical obligations that each domain generates. Chapter 5 (learning systems) is the most accessible entry point to Part Two; Chapter 7 (physical field systems) is the most technically demanding.

Part Three extracts a domain-independent diagnostic methodology from the results of Parts One and Two. The five diagnostic questions, the failure taxonomy, the alignment design principle, and the formal audit procedure provide the practical tools used throughout the remainder of the book.

Part Four applies the framework to domains that do not belong to the verified classes of Part Two: individual memory, institutional records, and cosmology. These chapters demonstrate the framework's research-program-generating function. The conclusions are frequently conditional, and open empirical obligations appear throughout. This is not a weakness—it is the framework operating correctly in domains where the evidence is incomplete.

Part Five develops the philosophical synthesis. It examines why the past matters less at higher interaction orders, what projection determines about which distinctions survive compression, and what the geometry of alternative histories reveals about the conditions for knowledge, explanation, and memory.

Recommended Reading Paths

For readers interested in the mathematics and formal theory: Read Parts One, Two, and Five in sequence.

For readers interested in machine learning, semantic systems, or physical field modeling: Read Part One in full, then focus on the relevant chapter in Part Two. Chapters 8 and 9 in Part Three develop the failure analysis and alignment principles that extend the Part Two results.

For readers interested in cognition or memory: Read Chapter 1, Chapter 10, and Chapter 11. Chapters 2–4 and Part Five provide formal and philosophical context.

For readers interested in institutions or historical records: Follow the same path as above, substituting Chapter 12. The failure-mode analysis of Chapter 8 maps particularly directly onto institutional phenomena.

For readers interested in cosmology or philosophy of physics: Read Chapter 1, Chapter 4, Chapter 7, and Chapter 13. Chapter 14 connects stability to the philosophy of explanation.

For readers interested primarily in the framework as a methodology: Begin with Chapter 10, which provides the complete five-question audit procedure, then use earlier chapters as references when formal justification is needed.

Three Distinctions to Keep in Mind

A **verified result** is a theorem, proof, or empirically supported claim established within a specified class of systems. Verified results appear primarily in Parts One and Two.

A **conditional result** is a statement whose validity depends on assumptions that remain empirically or mathematically open. Conditional results appear throughout Parts Three, Four, and Five.

An **open empirical obligation** is a measurement, simulation, proof, or experiment whose completion would resolve a conditional result. Open obligations are outputs of the framework: precise scientific questions generated by applying it to a domain where evidence is currently insufficient.

How to Use This Book

The chapters that follow move from mathematics to application and from application to philosophy. The order is deliberate. Each layer provides concepts that become precise only after the previous layer has been developed.

Readers interested in the destination may begin almost anywhere.

Readers interested in the foundations should begin at the beginning.

Acknowledgments

No work of this scope is created in isolation.

The framework developed in this book emerged from the intersection of several intellectual traditions that long predate it. The mathematical foundations draw from approximation theory, information theory, topology, dynamical systems, and the theory of generalized inverses. The conceptual foundations draw from machine learning, cognitive science, philosophy of science, legal theory, and cosmology. Although the framework assembles these elements into a single structure, its components were developed by many researchers across many disciplines over many decades. I am indebted to all of them, whether or not they appear in the bibliography that follows.

Particular gratitude is owed to the authors whose work provided the starting points for several of the framework's central ideas: the learning-theoretic reconstruction perspective developed by Gunn; influence-function methods introduced by Koh and Liang; the predictive-processing and free-energy traditions associated with Friston and others; the traditions of approximation theory and dynamical systems that supplied much of the mathematical vocabulary; and the broader literature on memory, institutional persistence, and cosmological inference that motivated the framework's applications.

The book also benefited from countless conversations, critiques, objections, and requests for clarification. Many of the ideas presented here became sharper because they survived disagreement. Errors that remain are entirely my own.

I am grateful as well to the communities of independent researchers, open-source contributors, writers, and educators whose willingness to share ideas publicly makes projects of this kind possible. A work that attempts to connect disciplines depends upon the existence of people willing to think beyond disciplinary boundaries.

Finally, I wish to acknowledge a more general debt. Every attempt to understand the past relies on records preserved by others: observations, experiments, archives, books, data sets, and memories that survived long enough to become part of the present. This book studies the conditions under which history remains recoverable after compression. It exists only because so much history was preserved before it.

The framework developed in these pages is offered not as a conclusion but as an invitation. The questions it raises are larger than the answers it provides. Many of its claims remain conditional. Many of its empirical obligations remain open. If the framework proves useful, it will do so because others extend it, challenge it, test it, and discover where it fails.

That is how knowledge advances. Not by preserving everything, but by preserving

enough.

Contents

Part I

Foundations

Chapter 1

Historical Dependence and Compression

1.1 The Problem of Incomplete Histories

Every record of the past is incomplete. This is not a limitation of current technology. It is the normal condition of any system that must act in the present on the basis of a history too large to store.

A machine-learning model is trained on millions of examples. When training ends, the parameters are stored; the training trajectory is not. A person accumulates experiences across a lifetime, but memory is a compression of that history, not a recording. A physical system evolves through microscopic configurations that its present macrostate does not encode.

In each case, a complex history has been replaced by a compressed representation, and questions persist about the history that was lost. What would the model have predicted had this training example been excluded? What would the person believe had certain experiences occurred differently? Which physical trajectories remain accessible from the present state?

1.2 Histories as Trajectories

A **trajectory space** $\mathcal{X}$ is the space of all possible histories of a system. A trajectory $x \in \mathcal{X}$ is a complete answer to the question: what actually happened? The trajectory space is therefore the space of alternative histories—the collection of all possible ways the system could have evolved.

Three trajectory spaces appear throughout this book: training histories for machine-learning models (the first verified domain), semantic inference trajectories (the second), and field configuration trajectories (the third).

1.3 Compression as Projection

A **compression operator** is a map

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

from the trajectory space to a representation space $\mathcal{M}$. It partitions $\mathcal{X}$ into equivalence classes, called **fibers**:

$$[x]_\pi = \pi^{-1}(\pi(x)) = \{x' \in \mathcal{X} : \pi(x') = \pi(x)\}.$$

The fiber $[x]_\pi$ is the set of all alternative histories the compression cannot distinguish from x. Two trajectories in the same fiber are indistinguishable given the compressed representation alone. The quotient space $\mathcal{X}/{\sim_\pi}$ is the "remembered world"—the world of recoverable inquiry.

1.4 Queries and Counterfactual Questions

A **query** is a function $Q : \mathcal{X} \to \mathbb{R}$. A **query family** $\mathcal{Q}$ is a collection of related queries. Queries that ask about relationships between the actual trajectory and hypothetical alternatives are **counterfactual queries**. They are the central objects of this book.

The query-relativity of reconstruction is a recurring theme: stability and sufficiency are always defined relative to a specific query family. There is no such thing as a history that is compressible in the abstract.

1.5 The Reconstruction Problem

The three-space structure is:

$$\mathcal{X} \xrightarrow{\pi} \mathcal{M}, \qquad Q : \mathcal{X} \to \mathbb{R}.$$

The reconstruction problem asks: when does there exist a map $R : \mathcal{M} \to \mathbb{R}$ such that $R(\pi(x)) \approx Q(x)$ for all x? Chapter 2 introduces the framework for answering this question precisely.

Chapter 2

The Fundamental Error Decomposition

2.1 Why Reconstruction Fails

Every failure of counterfactual reconstruction from a compressed representation can be traced to one or more of three independent sources. A model's sketch captures only low-order derivative interactions; the query depends on high-order ones—compressibility failure. A compression discards syntactic structure; the query requires it—sufficiency failure. An accessibility skeleton is stored to depth three; the query requires depth seven—approximation failure, or insufficient compression depth.

2.2 The Three Error Sources

Compressibility failure. The query's sensitivity to the history does not decay with interaction order. No finite compression can capture the history's full influence on the query.

Sufficiency failure. The compression discards information the query requires. The influence is in principle compressible, but the compression preserves the wrong structure.

Approximation failure. The query cannot be represented accurately at the chosen truncation order, even when the history is compressible and the compression is sufficient.

These three sources are independent: any one can occur while the others are absent.

2.3 The Fundamental Error Decomposition

Definition 2.1 (Fundamental Error Decomposition)**.** The total reconstruction error satisfies

$$E_{\text{total}} = E_c + E_s + E_a$$

where:

- $E_c = C_Q \sum_{r>d} \sigma_r$ is the **compressibility error**,
- $E_s = \varepsilon_1$ is the **sufficiency error**,
- $E_a = \varepsilon_0$ is the **approximation error**.

Each term has an independent source and an independent remedy.

2.4 Three-Space Structure

The reconstruction problem operates across three distinct spaces. The **trajectory space** $\mathcal{X}$ is the ontological space of all possible histories. The **representation space** $\mathcal{M}$ is the epistemic space of compressed states. The **query space** $\mathbb{R}$ is the pragmatic space of answers.

2.5 Pure Failure Modes and Alignment

The three pure failure modes—stability, sufficiency, and approximation—correspond to the three error terms dominating in isolation. In practice, all three contribute simultaneously, and diagnosis requires identifying which dominates.

An **aligned** compression is designed so that a single depth parameter d simultaneously controls all three terms. This eliminates the three architectural failure modes by construction. Alignment is defined precisely in Chapter 4 and is the central design principle of Chapter 9.

Chapter 3

Interaction Decompositions and Stability

3.1 How Queries Depend on Histories

A query $Q : \mathcal{X} \to \mathbb{R}$ depends on the trajectory x through its history. The central question is whether that dependence is concentrated at low *interaction orders*—whether the influence of increasingly complex historical interactions decays. When it does, the history is compressible for this query.

3.2 Interaction Decompositions

Fix a reference trajectory $x_0 \in \mathcal{X}$. An **interaction decomposition** expresses each query as

$$Q(x) = \sum_{r=0}^{\infty} I_r^{Q,x_0}(x)$$

where $I_r^{Q,x_0}(x)$ measures the contribution of r-th order interactions between x and x_0.

Definition 3.1 (Stable Interaction Decomposition)**.** The decomposition is **stable** with stability sequence $\{\sigma_r\}$ if

$$|I_r^{Q,x_0}(x)| \leq C_Q \cdot \sigma_r$$

for all $Q \in \mathcal{Q}$, all $x \in \mathcal{X}$, and all $r \geq 0$.

The three verified domains instantiate interaction order differently: polynomial degree r for learning systems, branch depth r for semantic systems, and transition depth r for physical field systems.

3.3 Stability Sequences and Tail Bounds

The crucial property of $\{\sigma_r\}$ is its tail behavior. If $\sum_{r>d} \sigma_r \to 0$ as $d \to \infty$, then truncating at order d yields controlled compressibility error. Under exponential stability $\sigma_r \leq Ce^{-cr}$:

$$\sum_{r>d} \sigma_r \leq \frac{Ce^{-cd}}{1 - e^{-c}}.$$

3.4 Why Stability Holds

Three mechanisms produce stability in the verified domains:

- **Derivative decay** in learning systems: higher-order partial derivatives of the training-to-output map decay exponentially.
- **Branching attenuation** in semantic systems: admissibility constraints suppress admissible trajectory volume geometrically.
- **Accessibility contraction** in physical field systems: accessible configuration volume contracts under admissible evolution.

The framework requires only that some mechanism produces the required decay, not any specific mechanism.

3.5 When Stability Fails

When $\{\sigma_r\}$ is not summable, E_c does not decrease with d. This is a structural property of the trajectory space for the query family—not a computational limitation that a better algorithm could overcome. Stability failure is the first and most fundamental diagnostic question.

3.6 The Stability Hypothesis as a Scientific Claim

In the learning and semantic domains, stability is proved from structural properties. In physical field systems, stability is an empirical hypothesis. The framework specifies what the hypothesis claims, what evidence would confirm it, and what each possible outcome implies.

Chapter 4

Stability, Sufficiency, and the Reconstruction Theorem

4.1 Query Sufficiency

Definition 4.1 (Query Sufficiency)**.** A compression $\pi : \mathcal{X} \to \mathcal{M}$ is (d, ε)-**sufficient** for a query family $\mathcal{Q}$ if for every $Q \in \mathcal{Q}$ there exists a map $F_d : \mathcal{M} \to \mathbb{R}$ such that

$$|Q_d(x) - F_d(\pi(x))| \leq \varepsilon$$

uniformly on $\mathcal{X}$, where Q_d is the degree-d approximation of Q.

Sufficiency is a relational property of the compression, the query family, and the truncation order. A compression sufficient for one query family may be insufficient for another.

4.2 The Reconstruction Theorem

Theorem 4.2 (Counterfactual Reconstruction)**.** *Let $\pi : \mathcal{X} \to \mathcal{M}$ be a compression and $\mathcal{Q}$ a query family. Assume:*

1. *Each $Q \in \mathcal{Q}$ admits a degree-d approximation Q_d with $|Q - Q_d| \leq \varepsilon_0$.*
2. *Q admits a stable interaction decomposition with stability sequence $\{\sigma_r\}$.*
3. *$\sum_{r>d} \sigma_r < \infty$.*
4. *π is (d, ε_1)-sufficient for $\mathcal{Q}$.*

Then the reconstruction operator $\pi_Q^+ := F_d$ satisfies

$$|Q(x) - (\pi_Q^+ \circ \pi)(x)| \leq C_Q \sum_{r>d} \sigma_r + \varepsilon_1 + \varepsilon_0 = E_c + E_s + E_a.$$

Under exponential stability, E_c decays exponentially in d.

The proof is an error-bookkeeping argument: split $Q = Q_d + R_d$ (approximation), apply stability to bound $|Q - Q_d|$ (compressibility), apply sufficiency to bound $|Q_d - F_d \circ \pi|$, then apply the triangle inequality.

4.3 Alignment

Definition 4.3 (Aligned Compression)**.** A compression π is **aligned** with stability sequence $\{\sigma_r\}$ and query family $\mathcal{Q}$ if, when $d_\pi = d_Q = d$:

1. π is $(d, 0)$-sufficient for $\mathcal{Q}_d$ (exact sufficiency),
2. $E_a = O(\sum_{r>d} \sigma_r)$ (inherited approximation decay), and
3. $E_{\text{total}} = O(\sum_{r>d} \sigma_r)$ (unified convergence).

An aligned compression forgets exactly what stability says can be forgotten and preserves exactly what stability says matters. This eliminates all three architectural failure modes simultaneously.

4.4 The Five Diagnostic Questions

The Reconstruction Theorem generates a complete diagnostic framework. For any system claiming counterfactual reconstruction:

1. What is the interaction decomposition, and is it exact?
2. Is the stability sequence summable?
3. Is the compression (d, ε_1)-sufficient for the query family?
4. Are d_π and d_Q aligned?
5. Does the equivalence relation avoid degeneracy?

A system that answers all five affirmatively, with quantitative bounds, has established its reconstruction claim rigorously. Part Two instantiates this framework in three verified domains.

Part II

Verified Domains

Chapter 5

Learning Systems and Algorithmic Histories

5.1 The Reconstruction Problem in Learning Systems

A machine-learning model is trained on a dataset of examples. When training ends, the parameters are fixed and the training history is discarded. The counterfactual question: what would the model have predicted had a particular training example been excluded?

Retraining large models for every such question is computationally prohibitive. The question is whether the model's current parameters contain enough information to answer such counterfactual questions without retraining.

5.2 Interaction Decomposition: Diagnostic Question 1

Let $\mathcal{X} = \mathbb{R}^n$, where x_i is a downweight on the i-th training example. The query is $Q(x) = g(f(x))$, the composition of the training map f with a measurement g. The reference is $x_0 = \mathbf{1}$ (full unweighted training). The Taylor expansion around x_0 provides the interaction decomposition:

$$Q(x) = \sum_{r=0}^{\infty} T_r(x - x_0),$$

where T_r is the degree-r homogeneous component. This decomposition is exact wherever Q is smooth. **Diagnostic Question 1: Affirmative.**

5.3 Stability: Diagnostic Question 2

Definition 5.1 (Gunn Stability Gunn [2026])**.** There exist $C, c > 0$ such that $\|T_r\| \leq Ce^{-cr}$ for all $r \geq 0$.

This says higher-order interactions among training examples decay exponentially. Setting $\sigma_r = Ce^{-cr}$ gives a summable stability sequence. **Diagnostic Question 2: Affirmative under Gunn's stability assumption** (empirically supported for small models; open for frontier scale).

5.4 Compression: Diagnostic Question 3

Gunn's polynomial sketch Gunn [2026] evaluates $Q(x_0 + \zeta\psi)$ for random complex directions ψ, extracting random projections of the derivative tensors via forward-mode automatic differentiation over truncated polynomial rings. The sketch achieves $(d, O(\varepsilon))$-sufficiency with sketch dimension $O(\log(1/\delta)/\varepsilon^2)$. **Diagnostic Question 3: Affirmative.**

5.5 Reconstruction and Alignment

By Theorem 4.2:

$$E_{\text{total}} \leq \frac{Ce^{-cd}}{1 - e^{-c}} + O(\varepsilon) + \varepsilon_0.$$

The sketch achieves **alignment by coincidence**: designed for computational efficiency, it happens to preserve derivative structure calibrated to the stability sequence.

5.6 Diagnostic Audit

Question	Status	Evidence
Decomposition exact?	Yes	Taylor series
Stability summable?	Yes (Gunn)	$\|T_r\| \leq Ce^{-cr}$, small-model tests
Compression sufficient?	Yes	Symmetric-subspace identity
Depths aligned?	Yes	$d_\pi = d_Q = d$
Degeneracy avoided?	Yes	Well-separated sketch classes

5.7 Lesson

Reconstruction works not because the training history is stored, but because the query depends on it in a compressible way. The stability assumption is the load-bearing hypothesis. Open obligation: verify exponential derivative decay at frontier-model scale.

Chapter 6

Semantic Systems and Constraint Grammars

6.1 The Reconstruction Problem in Semantic Systems

A reasoning system moves through a space of beliefs and inferences governed by admissibility constraints. Its current state is the result of a trajectory through that space that cannot be fully recovered from the state alone. Counterfactual questions: what would the system conclude had a premise been changed? Had a constraint been relaxed?

6.2 Semantic Trajectory Spaces

A **semantic state space** S is a finite set of admissible configurations. A **constraint grammar** $G = (S, R, C)$ specifies a transition relation $R \subseteq S \times S$ and admissibility constraints C. The **admissible successor operator** is

$$\mathrm{Succ}_C(s) = \{s' \in S : (s, s') \in R \text{ and } C \text{ permits } (s, s')\}.$$

The branching factor $B = \max_s |\mathrm{Succ}_C(s)|$. A **semantic trajectory** $(s_0, s_1, \ldots, s_T)$ with $s_{t+1} \in \mathrm{Succ}_C(s_t)$ forms the semantic trajectory space X_{sem}.

6.3 Branch-Depth Interaction Decomposition: Diagnostic Question 1

For query Q and reference x_0, define $Q_r(x)$ as the restriction to trajectories diverging from x_0 at depth $\leq r$, with $Q_{-1} \equiv 0$. The **branch-depth interaction term** is $I_r^{Q,x_0}(x) = Q_r(x) - Q_{r-1}(x)$. The decomposition $Q = \sum_r I_r^{Q,x_0}$ is exact by telescoping. **Diagnostic Question 1: Affirmative.**

6.4 Semantic Stability: Diagnostic Question 2

A grammar is α-**attenuating** if $N_r \leq (B\alpha)^r$ for all r, where N_r is the number of admissible trajectories diverging from x_0 at depth r.

Theorem 6.1 (Semantic Stability). *For an α-attenuating grammar with $B\alpha < 1$ and queries bounded by M:*

$$|I_r^{Q,x_0}(x)| \leq M(B\alpha)^r, \qquad \sum_{r>d} \sigma_r \leq \frac{M(B\alpha)^d}{1 - B\alpha}.$$

Diagnostic Question 2: Affirmative for α-attenuating grammars.

6.5 CLIO Compression: Diagnostic Question 3

The **CLIO operator** π_{CLIO} stores the branch tree $T_{d_\pi}(x)$ through depth d_π. Two trajectories are d_π-equivalent iff their branch trees agree. For $Q \in \mathcal{Q}_{d_Q}$ (depth-d_Q queries) with $d_\pi \geq d_Q$: $\varepsilon_1 = 0$ exactly. **Diagnostic Question 3: Affirmative, exactly.**

6.6 The Alignment Theorem

Theorem 6.2 (Stability–Sufficiency Alignment). *Under Theorem 6.1 with $d_\pi = d_Q = d$:*

$$E_s = 0, \quad E_c \leq \frac{M(B\alpha)^d}{1 - B\alpha}, \quad E_a \leq \frac{M(B\alpha)^d}{1 - B\alpha}, \quad E_{\text{total}} = O((B\alpha)^d).$$

CLIO achieves **alignment by design**: the branch-tree equivalence relation is explicitly calibrated to the stability structure. Sufficiency is not an independent obstruction.

6.7 BNF and Specification Independence

BNF grammars provide implementation-independent specifications of semantic trajectory spaces. Attenuation is a property of the grammar's production rules, verifiable before any implementation is built. This mirrors the framework's emphasis on specification-level diagnostics.

6.8 Diagnostic Audit

Question	Status	Evidence
Decomposition exact?	Yes	Telescoping
Stability summable?	Yes (verified class)	$B\alpha < 1$
Compression sufficient?	Yes, $\varepsilon_1 = 0$	Branch-tree preservation
Depths aligned?	Yes	Alignment theorem
Degeneracy avoided?	Yes (verified class)	Tree equivalence

Lesson: Stability is discovered; alignment can be engineered.

Chapter 7

Physical Fields and Accessibility Dynamics

7.1 The Central Asymmetry

A semantic system can be designed to attenuate. A physical system cannot. Its dynamics are whatever the field equations make them. This is the fundamental asymmetry between Chapters 6 and 7.

7.2 Field Grammars and Accessibility Trajectories

A **field configuration** is a tuple $F = (\Phi, \mathbf{v}, S)$ of fields on a domain. A **field grammar** $\mathcal{G} = (\mathcal{F}, \mathcal{T}, \mathcal{C})$ specifies the configuration space $\mathcal{F}$, transition operators $\mathcal{T}$, and admissibility constraints $\mathcal{C}$. An **accessibility trajectory** $(F_0, F_1, \ldots, F_T)$ has $F_{t+1} \in \mathrm{Succ}_{\mathcal{C}}(F_t)$.

7.3 Transition-Depth Interaction Decomposition: Diagnostic Question 1

For query Q and reference γ_0, define $A_r^{Q,\gamma_0}(\gamma) = Q_r(\gamma) - Q_{r-1}(\gamma)$ by the same telescoping construction as Chapter 6. The decomposition $Q = \sum_r A_r^{Q,\gamma_0}$ is exact by finite depth. **Diagnostic Question 1: Affirmative, unconditional.**

7.4 Accessibility Stability: Diagnostic Question 2

Let μ be a measure on $\mathcal{F}$ and $\mathcal{A}_r(F) = \{F' : F' \text{ reachable in exactly } r \text{ transitions}\}$.

Definition 7.1 (α-Contracting Field Grammar)**.** $\mathcal{G}$ is **α-contracting** if $\mu(\mathcal{A}_r(F)) \leq \alpha^r \mu(\mathcal{A}_0(F_0^{\mathrm{ref}}))$ for all F, r.

Theorem 7.2 (Accessibility Stability). *For an α-contracting grammar with $B_{\text{eff}} \cdot \alpha < 1$ and L-Lipschitz queries:*

$$|A_r^{Q,\gamma_0}| \leq L \cdot r \cdot (B_{\text{eff}}\alpha)^r, \qquad \sum_{r>d} \sigma_r = O(d \cdot (B_{\text{eff}}\alpha)^d).$$

Diagnostic Question 2: Conditional on α-contraction.

7.5 Accessibility Skeleton Compression: Diagnostic Question 3

The **accessibility skeleton** $\mathcal{S}_{d_\pi}(F) = \bigcup_{r=0}^{d_\pi} \mathcal{A}_r(F)$ stores all reachable configurations through depth d_π. Two trajectories are d_π-equivalent iff their skeletons agree μ-almost everywhere. For depth-d_Q accessibility queries with $d_\pi \geq d_Q$: $\varepsilon_1 = 0$ exactly (μ-a.e.). **Diagnostic Question 3: Affirmative, unconditional.**

7.6 Conditional Alignment and Three Outcomes

Under α-contraction with $d_\pi = d_Q = d$: $E_{\text{total}} = O(d(B_{\text{eff}}\alpha)^d)$.

The **Accessibility Relaxation Hypothesis** conjectures that RSVP field dynamics produce α-contracting grammars through systematic decrease of the accessibility field S. Three outcomes:

Outcome A (exponential contraction): Stability holds; full alignment achieved.

Outcome B (polynomial contraction): Polynomial stability; alignment partial; reconstruction possible but slow.

Outcome C (no contraction): Stability fails; E_c non-convergent; structural impossibility.

7.7 Diagnostic Audit

Question	Status	Evidence
Decomposition exact?	Yes, unconditional	Telescoping
Stability summable?	Conditional (Outcome A/B)	Open empirical question
Compression sufficient?	Yes, unconditional	Skeleton preservation
Depths aligned?	Conditional (Outcome A/B)	Requires contraction
Degeneracy avoided?	Requires verification	Measure-theoretic

Lesson: Alignment can be engineered, but stability cannot be assumed.

Part III

Diagnostics

Chapter 8

The Five Failure Modes

8.1 Failure as Information

A reconstruction failure is not the absence of an explanation. It is itself an explanation. The framework classifies failures precisely—distinguishing architectural mistakes from structural limits of the trajectory space.

The key distinction: **architectural failures** arise from design choices that can be improved. **Structural failures** arise from properties of the trajectory space that cannot be engineered away.

8.2 Stability Failure

Stability failure occurs when the stability sequence $\{\sigma_r\}$ is not summable. The compressibility error $E_c = C_Q \sum_{r>d} \sigma_r$ does not converge to zero at any finite truncation order. This is a property of the trajectory space for the query family—no compression, however well-designed, can remedy it.

Examples. Learning: derivative tensors fail to decay. Semantic: $B\alpha \geq 1$. Physical: Outcome C. Memory: trauma-related queries may exhibit this pattern.

Classification: Structural.

8.3 Sufficiency Failure

Sufficiency failure occurs when $d_\pi < d_Q$: the compression retains less branch structure than the query requires. The reconstruction operator F_{d_Q} cannot be well-defined because the compression collapses distinct trajectories that the query distinguishes.

Remedy: Increase d_π until $d_\pi \geq d_Q$. Classification: Architectural.

8.4 Approximation Failure

Approximation failure occurs when the query has significant interaction contributions at depths far exceeding any tractable truncation order. Even with perfect stability and sufficiency, E_a remains large.

Remedy: Increase d_Q or switch to a better-adapted approximation family. Classification: Architectural.

8.5 Misalignment Failure

Misalignment failure occurs when the compression preserves structure calibrated to a different mechanism than the stability sequence respects. Three forms: depth misalignment ($d_\pi < d_Q$), over-compression ($d_\pi > d_Q$), and structural misalignment (equivalence relation misoriented relative to stable directions).

Remedy: Redesign the compression around the stability mechanism. Classification: Architectural.

8.6 Degeneracy Failure

Degeneracy failure is introduced by measure-theoretic compressions and has no analogue in discrete settings.

Type I (collapsing). The compression identifies trajectories the query family treats as distinct. Information is destroyed. Examples: false memories, historical revisionism.

Type II (conditioning). The compression formally achieves $\varepsilon_1 = 0$ but the reconstruction operator is ill-conditioned: small perturbations in query-distinguishable trajectories produce large reconstruction output changes. Examples: accessibility skeleton degeneracy, dark-sector parameter degeneracy in cosmology.

Remedy: Enrich the equivalence relation or measure structure. Classification: Structural-Architectural boundary.

8.7 Architectural Versus Structural Failures

Failure Mode	Type	Remedy
Stability failure	Structural	None at compression level
Degeneracy (Type I/II)	Structural-Arch. boundary	Refine equivalence relation
Sufficiency failure	Architectural	Redesign compression
Approximation failure	Architectural	Better approximation family
Misalignment failure	Architectural	Calibrate to stability mechanism

A reconstruction failure is not the absence of an explanation. It is itself an explanation.

Chapter 9

Alignment as a Design Principle

9.1 From Failure Classification to Failure Prevention

Chapter 8 classified reconstruction failures. The natural follow-up: can compression architectures be built so that the architectural failures never arise?

Definition 9.1 (Aligned Compression Architecture)**.** A compression π is **aligned** with stability sequence $\{\sigma_r\}$ and query family $\mathcal{Q}$ if, with $d_\pi = d_Q = d$:

1. π is $(d, 0)$-sufficient for $\mathcal{Q}_d$.
2. $E_a = O(\sum_{r>d} \sigma_r)$.
3. $E_{\text{total}} = O(\sum_{r>d} \sigma_r)$.

Alignment transfers complexity from query time to design time. An aligned architecture performs the difficult work before the query arrives.

9.2 Three Notions of Alignment

Alignment by coincidence (learning systems). The polynomial sketch was designed for computational efficiency. It achieves alignment because the symmetric-subspace identity happens to preserve derivative structure calibrated to Gunn stability.

Alignment by design (semantic systems). CLIO was explicitly designed so that branch-tree preservation matches the stability structure of α-attenuating grammars. Sufficiency is exact by construction.

Alignment under uncertainty (physical fields). Accessibility skeleton compression is designed to align with hypothesized α-contraction. Sufficiency is guaranteed unconditionally; alignment of E_c and E_a is conditional on the Accessibility Relaxation Hypothesis.

9.3 The Economics of Alignment

An unaligned compression forces each query to pay its own complexity cost. An aligned compression pre-pays the complexity cost once, during construction. The larger the query family and the more frequently queries are asked, the more valuable this investment becomes.

9.4 The Design Principle

Identify the stability mechanism of the trajectory space. Calibrate the compression's equivalence relation to the equivalence classes that the stability mechanism respects. The query approximation family should be defined by dependence on the same equivalence classes. When the stability mechanism is known, this produces alignment by design. When hypothesized, it produces alignment under uncertainty.

An aligned compression architecture is one that performs the difficult work before the query arrives.

Chapter 10

The Five Diagnostic Questions as a Scientific Instrument

10.1 From Theorems to Audits

The five diagnostic questions of Chapter 4 operate at three levels. As *verification criteria*, they determine whether reconstruction is guaranteed. As a *scientific workflow*, they determine why reconstruction works or fails. As a *research program generator*, they determine what experiment, theorem, or measurement would resolve remaining uncertainty.

The purpose of the audit is not to certify reconstruction. It is to identify the minimum evidence required to justify a reconstruction claim.

10.2 The Five-Question Audit Procedure

Question 1: What is the interaction decomposition, and is it exact? Specify $\mathcal{X}$, x_0, the interaction order, and the terms I_r^{Q,x_0}. A positive answer establishes that the query's dependence on the history can be analyzed at all.

Question 2: Is the stability sequence summable? Specify $\{\sigma_r\}$ and the tail behavior. If not summable: *stability failure* (structural). Report as a property of the trajectory space; consider alternative decompositions.

Question 3: Is the compression sufficient? Specify π, d_π, and the sufficiency error ε_1. If large: *sufficiency failure* (architectural). Redesign the compression.

Question 4: Are depths aligned? Verify $d_\pi = d_Q$. If not: *misalignment failure*. Adjust depth parameters.

Question 5: Does the equivalence relation avoid degeneracy? Check whether $\sim_{d_\pi}$ collapses trajectories the query family distinguishes. If so: *degeneracy failure*. Enrich the equivalence relation or measure structure.

10.3 The Diagnostic Audit Report

A complete audit produces a nine-component report:

1. Trajectory Space Definition
2. Query Family Definition
3. Interaction Decomposition
4. Stability Analysis
5. Compression Architecture
6. Sufficiency Analysis
7. Alignment Analysis
8. Failure Classification
9. Open Empirical Obligations

Component 9 is the most important: it converts vague uncertainty into concrete scientific questions.

10.4 Research Programs from Audit Gaps

Open obligations from the three verified domains:

- **Learning systems:** Measure derivative tensor decay at frontier-model scale.
- **Semantic systems:** Measure attenuation coefficients for natural-language inference grammars.
- **Physical fields:** Determine whether accessible volume contracts exponentially, polynomially, or not at all under RSVP dynamics.

The framework transforms philosophical uncertainty about whether reconstruction is possible into empirical questions about specific measurable quantities.

10.5 Portability

The five-question audit is portable to any domain where counterfactual reconstruction from compressed representations is claimed or desired. It is modular, falsifiable, research-program-generating, and sensitive to epistemic levels (proved versus empirically supported versus hypothetical).

The purpose of Part Three has not been to prove that reconstruction succeeds. It has been to determine what evidence would be required to believe that it succeeds.

Part IV

Applications

Chapter 11

Cognition and Memory

11.1 What Kind of Object Is a Memory?

From the framework's perspective, memory is a compression operator on experiential trajectories. The neural mechanisms, storage architecture, and retrieval dynamics are implementation details. What matters is the compression: what equivalence classes it induces on experiential trajectories, what queries it supports, and whether it is aligned with the stability structure of experiential dependence.

Before asking whether memory is stable or sufficient, the framework must ask: what is the trajectory space?

11.2 Candidate Trajectory Spaces

Three candidates arise:

Episodic. A trajectory is a temporally ordered sequence of experiences. Counterfactual question: what would I believe had experience X been replaced by Y?

Inferential. A trajectory is a sequence of belief states and inferential transitions. Counterfactual question: what would I conclude had I inferred differently at some earlier stage?

Predictive. A trajectory is a sequence of action-perception loops with continuous prediction-error updates. Connects to the free-energy principle Friston [2010].

Diagnostic Question 1: Partially open. Multiple candidate interaction orders exist (temporal, associative, causal depth); none is established as canonical.

11.3 Stability: Diagnostic Question 2

The forgetting curve Ebbinghaus [1913] suggests $\sigma_r \sim r^{-k}$ under temporal interaction order—polynomial decay, consistent with a summable stability sequence. Interference effects complicate the picture; traumatic experiences may exhibit stability failure for specific query families.

Diagnostic Question 2: Conditional. Evidence suggests polynomial decay; exponential stability unestablished; trauma may produce stability failure.

11.4 Compression and Failure Modes

Memory induces an equivalence relation on trajectories. Candidate relations: semantic equivalence, narrative equivalence, emotional equivalence, predictive equivalence.

The failure-mode taxonomy maps cleanly onto recognized phenomena:

- **False memories** $\leftrightarrow$ Type I degeneracy
- **Amnesia** $\leftrightarrow$ sufficiency failure
- **Overgeneralization** $\leftrightarrow$ structural misalignment
- **Intrusive memories** $\leftrightarrow$ stability failure (trauma)

11.5 Diagnostic Audit and Open Obligations

Question	Status	Notes
Decomposition exact?	Partially open	No canonical interaction order
Stability summable?	Conditional	Polynomial decay suggested
Compression sufficient?	Conditionally yes	Equivalence relation unknown
Depths aligned?	Unknown	Depends on Q1, Q2
Degeneracy avoided?	Active concern	False memories suggest Type I

Open obligations: (1) determine canonical interaction order; (2) measure stability sequence; (3) characterize memory equivalence classes; (4) measure reconstruction conditioning; (5) identify stability structure of traumatic memory.

Chapter 12

Institutions and Historical Records

12.1 What Kind of Object Is an Institution?

An institution is a compression architecture for collective history. Courts compress legal history into precedent and doctrine. Governments compress political history into laws and regulations. Archives compress documentary history into retrievable records. The trajectory space—sequences of decisions, records, laws, and organizational transitions—is substantially more tractable than the experiential trajectory space of Chapter 11.

12.2 Institutional Trajectory Spaces

Several types arise: legal (cases, statutes, judicial interpretations), political (elections, legislative acts), organizational (hiring, strategy, culture), and scientific (experimental results, publications, citations). Each supports different counterfactual questions.

Diagnostic Question 1: Mostly affirmative. Precedent depth for legal systems mirrors branch depth from Chapter 6, providing a tractable interaction order. Legal citation networks Landes and Posner [1976] provide an empirical proxy.

12.3 Stability: Diagnostic Question 2

Institutional influence exhibits heavy-tailed decay. Constitutional provisions from centuries ago retain binding authority; Roman law concepts persist across modern jurisdictions. Scientific citation half-lives de Solla Price [1965] are finite, suggesting polynomial decay. Religious and cultural institutions may exhibit stability failure under temporal interaction order.

Diagnostic Question 2: Conditional, polynomial stability likely.

12.4 Institutional Compressions: Diagnostic Question 3

Institutional compressions are explicit and inspectable: legal archives, accounting records, scientific publications, oral traditions. Sufficiency is high for contemporary institutional queries; conditional for historical counterfactual queries not anticipated during compression design.

12.5 Failure Modes and Institutions

- **Archive destruction** $\leftrightarrow$ sufficiency failure (irreversible)
- **Historical revisionism** $\leftrightarrow$ Type I degeneracy (compression collapses distinct trajectories)
- **Bureaucratic overgeneralization** $\leftrightarrow$ structural misalignment
- **Institutional lock-in** $\leftrightarrow$ stability failure for reform queries
- **Lost oral tradition** $\leftrightarrow$ compression-operator destruction

12.6 Diagnostic Audit and Open Obligations

Question	Status	Notes
Decomposition exact?	Mostly affirmative	Precedent depth tractable
Stability summable?	Conditional	Polynomial likely; varies by domain
Compression sufficient?	Affirmative (designed queries)	Conditional (historical queries)
Depths aligned?	Variable	Legal systems better than others
Degeneracy avoided?	Active concern	Revisionism demonstrates Type I

Open obligations: (1) measure institutional influence decay; (2) characterize precedential equivalence; (3) measure archival sufficiency for counterfactual queries; (4) develop alignment-based archival design principles; (5) identify boundary between polynomial and non-summable stability for cultural institutions David [1985], Arthur [1994].

Chapter 13

Cosmology and the Universe as a Compression

13.1 What Kind of Object Is a Cosmological Model?

A cosmological model is a compression architecture for physical history. The observable universe contains only a fraction of the information in its past trajectory. The CMB encodes conditions at recombination but not the precise configuration of every field mode. Large-scale structure encodes gravitational amplification of primordial perturbations but not every fluid element's trajectory.

A cosmological theory specifies what information is preserved in the present state, what is discarded, and which historical queries can in principle be answered.

13.2 Cosmological Trajectory Spaces

Four candidates: (1) spacetime-history (complete field histories), (2) cosmological-state (large-scale parameter sequences), (3) accessibility (RSVP field configurations), and (4) observer-conditioned (histories relative to the observer's causal past). The canonical trajectory space is not settled across cosmological theories.

Diagnostic Question 1: Affirmative within specific formalisms; contested across them.

13.3 Stability: Diagnostic Question 2

Λ**CDM.** Perturbation theory suggests a summable stability sequence for large-scale observables under temporal interaction order.

Inflation. The trans-Planckian problem Martin and Brandenberger [2001], Baumann [2009] is a stability question: if Planck-scale physics affects inflationary observables, the stability sequence is not summable under energy-scale interaction order.

RSVP. The Accessibility Relaxation Hypothesis conjectures $\mu(\mathcal{A}_r(F_0)) \leq Ce^{-cr}$ under RSVP dynamics. Three outcomes from Chapter 7 apply at cosmological scale.

Diagnostic Question 2: Conditional across all frameworks.

13.4 Cosmological Failure Modes

- **Trans-Planckian problem** $\leftrightarrow$ stability failure candidate
- **Cosmic variance** $\leftrightarrow$ sufficiency failure (imposed by causal structure)
- **Model underdetermination** $\leftrightarrow$ structural misalignment
- **Information loss behind horizons** $\leftrightarrow$ stability failure or sufficiency failure
- **Dark-sector parameter degeneracy** $\leftrightarrow$ Type II degeneracy

13.5 RSVP Audit

RSVP is audited as a candidate theory, not a confirmed one.

Question	Status	Notes
Decomposition exact?	Yes	Accessibility telescoping
Stability summable?	Conditional	Accessibility Relaxation Hypothesis
Compression sufficient?	Yes	Skeleton compression
Depths aligned?	Conditional	Requires contraction
Degeneracy avoided?	Open	Measure-theoretic

Open obligations: (1) derive accessibility contraction from RSVP field equations; (2) identify observable signatures distinguishing three outcomes; (3) test inflationary stability against trans-Planckian corrections Planck Collaboration [2020]; (4) develop compression-theoretic model distinguishability measures; (5) characterize dark-sector degeneracy structure.

13.6 The Universe as a Compression of Its Own History

If cosmological history is compressible (Outcome A), the present state contains recoverable information about nearby alternative histories. If it is not (Outcome C), the limits of cosmological knowledge are structural, not practical. Both possibilities have precise meanings in the framework's language. Both can in principle be distinguished by empirical investigation.

Part V

Geometry of Alternative Histories

Chapter 14

Why the Past Matters Less

14.1 A Recurring Question

Throughout this book a single question has reappeared in every domain: does historical influence decay with interaction order? In learning systems it appeared as derivative decay. In semantic systems as branch attenuation. In physical fields as accessibility contraction. In memory as the forgetting curve. In institutions as precedential half-lives. In cosmology as accessible volume decay.

This recurrence is not an accident of the framework's design. Stability is a structural property of historical systems themselves.

14.2 Three Faces of Stability

Stability as forgetting. In memory systems, stability appears as forgetting—the progressive reduction of sensitivity to remote experiential interactions. A world without forgetting would be a world where nothing could be compressed, summarized, or reconstructed.

Stability as abstraction. In semantic systems, higher-order branch interactions are absorbed into coarser conceptual structures. Abstraction is the cognitive manifestation of stability.

Stability as relaxation. In physical systems, accessible configuration volume contracts as field dynamics evolve. The system loses sensitivity to increasingly remote alternative initial conditions.

14.3 Stability as Historical Curvature

A stable trajectory space is one in which nearby alternative histories *converge under projection*: as truncation depth d increases, the distinctions among trajectories within the same fiber become progressively less significant for the query.

An unstable trajectory space is one in which nearby histories *diverge indefinitely*: no finite compression makes the fiber query-flat.

The Stability Principle: a trajectory space is stable for a query family when the projection of nearby trajectories onto query-answer space contracts with increasing interaction order.

14.4 The Three Historical Universes

Outcome A (exponential stability). History becomes rapidly irrelevant. Compression is efficient; reconstruction is tractable. Science is easy; institutions are reformable.

Outcome B (polynomial stability). History remains important for a long time. Compression is possible but costly; reconstruction requires large depth.

Outcome C (no stability). History never becomes negligible. Compression fails; reconstruction is impossible; explanation itself breaks down.

14.5 Stability and Explanation

Why do explanations work? The framework's answer: explanations work because most historical influence eventually becomes negligible. A good explanation is a compression sufficient for the query families relevant to understanding and prediction. In a world without stability, nothing could be summarized, modeled, or reconstructed. The possibility of explanation is a consequence of stability.

14.6 Query-Relativity of Stability

Stability is not intrinsic to the trajectory space. It is a relational property of the trajectory space, the query family, and the interaction decomposition. A physical system may be exponentially stable for thermodynamic queries and chaotic for microscopic queries. This query-relativity prevents treating "stable systems" as an intrinsic category.

Stability determines when compression is possible. Projection determines what is preserved.

Chapter 15

Projection and Representation

15.1 Why This Projection Rather Than Another?

Parts One through Four largely treated the compression operator π as given. Chapter 15 asks the prior question: why is one compression chosen rather than another?

Different compressions of the same trajectory space create different worlds of recoverable history. Two observers compressing the same $\mathcal{X}$ using different operators inherit different equivalence relations, different query families for which reconstruction is possible, and different failure modes.

15.2 Projection Creates Worlds

The compression $\pi : \mathcal{X} \to \mathcal{M}$ partitions $\mathcal{X}$ into fibers $\pi^{-1}(m)$. The recovered world is not $\mathcal{X}$ but the quotient $\mathcal{X}/{\sim_\pi}$—the world of equivalence classes, where some differences have been erased and others preserved.

Memory preserves semantic content while discarding perceptual detail. Archives preserve formal reasoning while discarding social context. Physical dynamics preserve macroscopic observables while discarding microscopic states. Scientific theories preserve invariant structure while discarding measurement noise. In each case the projection *creates* the world available to reconstruction.

15.3 Projection as Structured Forgetting

A representation is a pattern of preserved distinctions. Projections forget selectively: they collapse some distinctions while preserving others. What a representation says is determined not by what it contains but by what it *excludes*—which trajectories it rules out.

Every domain in this book exemplifies this: sketches preserve derivative structure; branch trees preserve admissible branch structure; skeletons preserve accessible future structure; memory preserves semantic-narrative structure; archives preserve doctrinal structure.

15.4 The Pseudoinverse Perspective

When π is many-to-one, exact inversion is impossible. The query-relative reconstruction operator $\pi_Q^+ : \mathcal{M} \to \mathbb{R}$ is a *query-relative pseudoinverse* Penrose [1955], Ben-Israel and Greville [2003]. It selects a canonical representative from each fiber and evaluates the query on it.

Reconstruction is always about the projection of the past, never the past directly. The accuracy of reconstruction depends not on choosing the correct representative but on whether all representatives in the fiber produce approximately the same query answer—which is precisely the stability condition.

15.5 Alignment as Geometric Property

An aligned projection collapses trajectories in *stable directions* (where σ_r decays rapidly) and preserves trajectories in *important directions* (where σ_r remains large at low r). Geometrically: the projection's null space is parallel to the stable directions.

Misalignment is geometrically transparent: sufficiency failure occurs when the projection collapses important directions; over-compression when it preserves stable directions unnecessarily.

15.6 Representation as Constraint

Representations are not primarily depictions. They are constraints: specifications of which trajectories remain possible. The significance of a representation is what it excludes, not what it contains. This constraint-based view explains why representations can be accurate without being complete, why multiple representations can be equally valid for the same query family, and why representation precedes inquiry.

15.7 Projection Before Knowledge

Before a system can know anything, it must project. Memory must compress before it can retrieve. Archives must be built before they can be consulted. Theories must model before they can predict. The world of recoverable inquiry is always defined before the inquiry begins. Knowledge is reconstruction performed inside a projection-induced quotient space.

Stability determines when compression is possible. Projection determines what is preserved. Together they define the geometry of recoverable pasts.

Chapter 16

The Geometry of Alternative Histories

16.1 What Is an Alternative History?

An alternative history is not merely a different sequence of events. It is a trajectory in $\mathcal{X}$ that occupies a specific position relative to the reference trajectory x_0, the query family $\mathcal{Q}$, the compression π, and the stability sequence $\{\sigma_r\}$. Its recoverability is a geometric property of its position in this structure.

16.2 The Three Spaces

The trajectory space $\mathcal{X}$ is the *ontological space* of all possible histories. The representation space $\mathcal{M}$ is the *epistemic space*. The query space $\mathbb{R}$ is the *pragmatic space*. These are connected by:

$$\mathcal{X} \xrightarrow{\pi} \mathcal{M}, \qquad Q : \mathcal{X} \to \mathbb{R}.$$

Reconstruction asks when there exists $R : \mathcal{M} \to \mathbb{R}$ making this diagram approximately commute.

16.3 Alternative Histories as Fibers

For each $m \in \mathcal{M}$, the **fiber**

$$\pi^{-1}(m) = \{x \in \mathcal{X} : \pi(x) = m\}$$

is the set of all trajectories the compression cannot distinguish from any member of the class. The actual trajectory x_0 lies in the fiber $\pi^{-1}(\pi(x_0))$; every other element of that fiber is an alternative history consistent with the compressed state.

Alternative histories are not hypothetical trajectories outside knowledge. They are trajectories inside the fiber of what current knowledge permits. Counterfactual reasoning is navigation within fibers.

16.4 Stability Shapes the Fiber

Without stability, fibers remain query-curved: different elements produce very different query answers. With stability, fibers become query-flat: the query answers of all trajectories within the same fiber converge as truncation depth d increases.

The Fiber Convergence Principle: stability does not eliminate alternative histories. It makes them query-convergent.

16.5 The Cone of Recoverability

Definition 16.1 (Cone of Recoverability)**.**

$$C_\varepsilon(x_0, \pi, \mathcal{Q}) = \{x \in \mathcal{X} : |Q(x) - \pi_Q^+(\pi(x))| \leq \varepsilon \text{ for all } Q \in \mathcal{Q}\}.$$

Under exponential stability (Outcome A): the cone is broad. Under polynomial stability (Outcome B): the cone is narrower. Under no stability (Outcome C): the cone collapses. The cone encodes the geometry of reconstruction for a given system.

16.6 The Geometry of Failure

- **Stability failure:** diverging fibers—the fiber is query-curved at every depth.
- **Sufficiency failure:** query-crossing fibers—the fiber spans multiple query contours.
- **Approximation failure:** fine-scale query curvature within the resolution of the compression.
- **Misalignment:** misoriented projection—the null space is perpendicular to the stable directions.
- **Degeneracy:** folded quotient—the fiber contains trajectories from widely separated regions of $\mathcal{X}$.

16.7 A World That Can Remember

For memory to exist, for explanation to exist, for science to exist, four geometric conditions must hold:

1. The trajectory space must be stable for the query family (fibers query-flat).
2. The projection must be sufficient (fibers do not cross query contours).
3. The projection must be aligned with the stability structure.
4. The equivalence relation must be well-conditioned for the query family.

When all four hold, the world remembers. When any fails, memory is impaired in a specific, geometrically characterizable way.

16.8 The Geometry of Recoverable Pasts

The trajectory space $\mathcal{X}$ contains all histories the world might have had. The stability sequence $\{\sigma_r\}$ determines how rapidly alternative histories converge in query-relevant directions. The compression π determines which distinctions survive. The fiber $\pi^{-1}(\pi(x_0))$ contains all alternative histories consistent with the present. The reconstruction operator π_Q^+ navigates the fiber. The cone $C_\varepsilon(x_0, \pi, \mathcal{Q})$ specifies which alternatives are recoverable.

The central question of this book was never whether the past can be reconstructed. It was what kind of world would make reconstruction possible. The answer is now visible. A world can remember its past when its histories contract, its projections preserve the right distinctions, and its alternative trajectories converge in the directions that matter. The geometry of alternative histories is therefore the geometry of recoverable memory itself.

The world remembers.

Appendix A

Notation, Glossary, and Framework Summary

A.1 Notation Table

Symbol	Meaning	First appears
$\mathcal{X}$	Trajectory space	Ch. 1
$\mathcal{M}$	Representation space	Ch. 1
$\mathbb{R}$	Query space	Ch. 1
$\pi : \mathcal{X} \to \mathcal{M}$	Compression (projection) operator	Ch. 1
$\sim_\pi$	Equivalence relation induced by π	Ch. 1
$[x]_\pi = \pi^{-1}(\pi(x))$	Fiber (equivalence class) of x	Ch. 1
$\pi^{-1}(m)$	Fiber of compressed state m	Ch. 16
$\mathcal{X}/\sim_\pi$	Quotient space	Ch. 15
$Q : \mathcal{X} \to \mathbb{R}$	Query function	Ch. 1
$\mathcal{Q}$	Query family	Ch. 1
Q_d	Degree-d approximation of Q	Ch. 2
$F_d : \mathcal{M} \to \mathbb{R}$	Reconstruction operator	Ch. 2
π_Q^+	Query-relative reconstruction operator	Ch. 2
π^+	Moore–Penrose pseudoinverse	Ch. 15
x_0, γ_0	Reference trajectory	Ch. 2
I_r^{Q,x_0}	r-th interaction term	Ch. 3
A_r^{Q,γ_0}	r-th accessibility interaction term	Ch. 7
$\{\sigma_r\}$	Stability sequence	Ch. 3
d_Q	Query truncation order	Ch. 4
d_π	Compression depth	Ch. 4
E_c	Compressibility error	Ch. 2
E_s	Sufficiency error	Ch. 2
E_a	Approximation error	Ch. 2
E_{total}	Total reconstruction error	Ch. 2
B	Branching factor (discrete)	Ch. 6
B_{eff}	Effective branching factor (continuous)	Ch. 7
α	Attenuation / contraction parameter	Ch. 6
μ	Measure on field configuration space	Ch. 7
$\mathcal{A}_r(F)$	Accessible r-neighborhood of F	Ch. 7
$\mathcal{S}_{d_\pi}(F)$	Accessibility skeleton	Ch. 7
$C_\varepsilon(x_0, \pi, \mathcal{Q})$	Cone of recoverability	Ch. 16
$G = (S, R, C)$	Semantic constraint grammar	Ch. 6
$\mathcal{G} = (\mathcal{F}, \mathcal{T}, \mathcal{C})$	Field grammar	Ch. 7
$X_{\text{sem}}, M_{\text{sem}}$	Semantic trajectory/representation spaces	Ch. 6

A.2 Key Definitions

Fiber. $\pi^{-1}(m) = \{x \in \mathcal{X} : \pi(x) = m\}$. The set of all alternative histories consistent with compressed state m. Counterfactual reasoning is navigation within fibers.

Query-flat fiber. A fiber in which all elements produce approximately the same query answers: $\sup_{x,x' \in \pi^{-1}(m)} |Q(x) - Q(x')| \leq \varepsilon$. The geometric consequence of stability.

Cone of recoverability. $C_\varepsilon(x_0, \pi, \mathcal{Q}) = \{x : |Q(x) - \pi_Q^+(\pi(x))| \leq \varepsilon\}$. Broad under exponential stability; narrow under polynomial; collapsed under no stability.

Aligned compression. A compression whose equivalence classes coincide with the directions the stability sequence renders negligible. Collapses stable directions; preserves important directions. Eliminates all architectural failure modes.

Type I degeneracy. Compression collapses trajectories the query family distinguishes. Information destroyed. Examples: false memories, revisionism.

Type II degeneracy. Formal sufficiency ($\varepsilon_1 = 0$) but ill-conditioned reconstruction. Examples: accessibility skeleton on continuous spaces, dark-sector parameter degeneracy.

A.3 Framework Summary: Eight Steps

1. Define the trajectory space $\mathcal{X}$ and reference x_0.
2. Define the interaction decomposition: specify I_r^{Q,x_0} and verify exactness.
3. Verify stability: bound $|I_r^{Q,x_0}| \leq C_Q \sigma_r$ and check $\sum_{r>d} \sigma_r \to 0$.
4. Define the compression π and depth d_π; identify the fiber structure.
5. Verify sufficiency: compute or bound ε_1.
6. Bound approximation error ε_0.
7. Check alignment: $d_\pi = d_Q$; all three error terms share the same decay rate.
8. Diagnose failure modes: which of the five failure types is active?

A.4 Domain Comparison

	Ch. 5 (Learning)	Ch. 6 (Semantic)	Ch. 7 (Physical)
Interaction order	Polynomial degree	Branch depth	Transition depth
Stability mechanism	Derivative decay	α-attenuation	α-contraction
Compression	Polynomial sketch	Branch tree	Accessibility skeleton
Sufficiency error	$O(\varepsilon)$	0 (exact)	0 (μ-a.e.)
Alignment type	By coincidence	By design	Under uncertainty
Verified class	MicroGPT (empirical)	$B\alpha < 1$	$B_{\text{eff}}\alpha < 1$
Open question	Frontier stability	Natural-language attenuation	RSVP contraction

Bibliography

W. Brian Arthur. *Increasing Returns and Path Dependence in the Economy*. University of Michigan Press, 1994.

Daniel Baumann. Inflation. *arXiv*, 2009. arXiv:0907.5424.

Adi Ben-Israel and Thomas N. E. Greville. *Generalized Inverses: Theory and Applications*. Springer, 2003.

Paul A. David. Clio and the economics of QWERTY. *American Economic Review*, 75:332–337, 1985.

Derek J. de Solla Price. Networks of scientific papers. *Science*, 149:510–515, 1965.

Hermann Ebbinghaus. *Memory: A Contribution to Experimental Psychology*. Teachers College, Columbia University, 1913.

Karl Friston. The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11:127–138, 2010.

Sam Gunn. How to sketch a learning algorithm. *arXiv*, 2026. arXiv:2604.07328.

William M. Landes and Richard A. Posner. Legal precedent: A theoretical and empirical analysis. *Journal of Law and Economics*, 19:249–307, 1976.

Jerome Martin and Robert H. Brandenberger. The trans-Planckian problem of inflationary cosmology. *Physical Review D*, 63:123501, 2001.

Roger Penrose. A generalized inverse for matrices. *Mathematical Proceedings of the Cambridge Philosophical Society*, 51:406–413, 1955.

Planck Collaboration. Planck 2018 results: Cosmological parameters. *Astronomy & Astrophysics*, 641:A6, 2020.

Index