Constraint as Ground

A Process Ontology of Cognition, Computation, and Cosmos

Flyxion

Independent Researcher

`https://github.com/standardgalactic`

2026

Prefatory Note: On the Form of This Text

This monograph does not proceed from a completed theoretical edifice toward illustrative applications. It proceeds in the opposite direction: from distributed conceptual material toward progressive topological stabilization. The reader should expect definitions to accumulate rather than appear fully formed, and should treat apparent repetition across sections as triangulation rather than redundancy.

The writing style is a mechanistic consequence of the theory rather than a stylistic preference. The core argument of Part V is that any text arguing for provenance preservation and reconstruction transparency while itself operating through aggressive compression and assumed priors would be internally inconsistent. The prefatory note, the proliferating appendices, the deliberate restating of assumptions at each new level of abstraction are not ornamental scholarly rituals. They are consequences of taking seriously the claim that meaning is reconstructed rather than transmitted.

Most of the conceptual work assembled here already existed as a large-scale distributed semantic field before this document imposed topological structure on it. Different companion papers foreground different regions of the same underlying geometry: cognition, cosmology, social systems, abstraction, entropy, compression, trajectory selection, admissibility, reconstruction, provenance. The apparent diversity of subjects conceals a relatively small set of deep structural operations being reapplied across domains and scales.

A note on formal notation. Definitions, theorems, propositions, and axioms are presented in numbered environments so that dependency relations can be tracked explicitly. Where a later argument depends on an earlier definition, the dependency is marked. Where a concept is imported from a companion document rather than derived here, the source is indicated. Terminology is introduced progressively and is not assumed stable across disciplinary boundaries.

Flyxion
Canada, 2026

Methodological Commitments

A recurring mathematical structure licenses transfer of operational insight only insofar as the transferred structure preserves measurable constraints, exclusions, and failure conditions.

This monograph tepflyxion2025collapse ranges across cosmology, cognition, computation, institutional analysis, and paleoarchaeology. That range creates a specific risk: that recurring mathematical structures will be taken as evidence of shared ontology rather than as evidence that the same formal tools apply usefully across different substrates. This chapter states the methodological commitments that guard against that risk.

The Analogy–Constraint Distinction

A recurring mathematical structure licenses transfer of operational insight only insofar as the transferred structure preserves *measurable constraints, exclusions, and failure conditions*. Shared formalism is not evidence of identical ontology. It is evidence that the same class of constraints operates at the relevant scale.

This principle is the methodological spine of the monograph. Whenever a formal framework developed in one domain is applied to another—when sheaf cohomology appears in both semantic analysis and social institution modeling, or when the RSVP entropy field appears in both cosmological and cognitive contexts—the transfer is licensed only by demonstrating that the relevant constraints, failure modes, and exclusion conditions are structurally analogous. The transfer is not licensed by aesthetic resonance, terminological similarity, or the intuition that the domains "feel" related.

The distinction is between *analogy* and *operational constraint transfer*. An analogy asserts resemblance without specifying which properties are preserved and which are not. An operational constraint transfer specifies exactly which constraints are shared, which are not, and what would falsify the claim that the shared constraints are sufficient to ground the transfer. Throughout this monograph, cross-domain applications are intended as constraint transfers rather than analogies.

Claim-Status Architecture

Because the monograph spans formally derived results, phenomenological interpretations, engineering heuristics, conjectures, and speculative extrapolations, it uses a consistent claim-status system to distinguish them. The five claim-status markers are:

[D] Definition. A formal stipulation establishing terminology or a formal object. Definitions are not true or false but more or less useful. They can be evaluated by whether they generate productive inferences.

[P] Phenomenological interpretation. An application of a formal structure to describe observable patterns without claiming that the formal structure is the underlying mechanism. Phenomenological interpretations can be accurate descriptions while remaining silent on mechanism.

[H] Heuristic principle. An operational guideline supported by evidence and useful in practice, but not derived from first principles and potentially wrong in edge cases. Heuristics should be tested against their domain of application.

[C] Conjecture. A claim that appears plausible given the available evidence and the internal logic of the framework, but has not been formally derived or empirically confirmed. Conjectures are productive when they generate testable predictions.

[S] Speculative extrapolation. A claim that goes beyond what the available evidence and formal derivations support, offered as a direction for investigation rather than an established result.

These markers appear in the text at the points where the claim-status of a statement shifts. They are not exhaustive annotations of every sentence: formally derived theorems and definitions carry their status implicitly. The markers appear where there is meaningful risk of status confusion—where a phenomenological interpretation might be read as a formal claim, or where a conjecture might be mistaken for an established result.

The most important risk is the silent transition from **[H]** to **[P]** to **[D]**: a heuristic becomes a phenomenological description becomes a formal definition without any single step being flagged as an escalation. The claim-status system exists to make those transitions visible.

Constraint Boundaries

Each major chapter ends with a *Constraint Boundary* section that states explicitly: what the framework in that chapter explains;

what it does not explain; what remains speculative; what would falsify the model; and which claims are phenomenological versus formal. These sections are intended as internal audits rather than summaries. They expose the limits of each local framework before the next chapter extends it.

The monograph consistently argues against opaque abstraction—against systems that compress away their dependency structure and present only terminal conclusions. The Constraint Boundary sections are the mechanism by which the monograph resists applying that pressure to itself.

On the RSVP Cosmological Material

One specific clarification is required at the outset regarding the RSVP cosmological framework. The entropic relaxation account of cosmological redshift and the admissibility geometry approach to structure formation are presented as **[C]** conjectures with explicit falsification criteria, not as established physics or replacements for the standard cosmological model. The redshift formula $z_{\text{RSVP}} \approx \int_0^L \kappa |\nabla S(\ell)| d\ell$ is a prediction of the RSVP field dynamics under specific coupling assumptions; it makes distinct predictions at $k \gtrsim 0.1\, h/\text{Mpc}$ and at the Silk damping scale that could in principle falsify the framework. Until those predictions are tested, the cosmological material functions as a geometric field formalism for representing structured evolution under entropy gradient, not as a replacement for established observational cosmology.

Contents

Part I

The Problem of Topology

> The difficult part is no longer generating isolated insights. The difficult part is imposing recoverable topology on a massive graph of partially overlapping trajectories.

Most of the conceptual work already exists as a distributed semantic field with diffuse boundaries. A monograph functions partly as a constraint geometry over that field: it establishes admissible traversal paths through the corpus so readers can progressively stabilize the conceptual dependencies rather than encountering them as isolated fragments.

Part I establishes why a unified process ontology requires a new traversal architecture rather than a new set of isolated claims, introduces the primitive vocabulary on which the remainder of the monograph is built, and argues that the form of the text is a consequence of its content.

Chapter 1

The Distributed Manifold Problem

> Different documents foreground different regions of the structure. The apparent diversity of subjects is masking a relatively small set of deep structural operations being reapplied across domains.

1.1 Why the Corpus Exists in the Form It Does

The body of work assembled across companion documents does not constitute a collection of unrelated essays, tools, experiments, games, cosmological proposals, audio systems, and philosophical notes. It constitutes a persistent attempt to construct a unified operational language for structure, cognition, computation, geometry, and social coordination. The repetition of certain motifs across wildly different domains is not accidental. The same underlying intuitions reappear because the projects are converging on a common claim: complex systems can often be understood as constrained transformations through spaces of admissible trajectories rather than as collections of isolated objects with intrinsic static identities (Flyxion, 2026a,b).

The problem this creates for a monograph is topological rather than creative. The material does not need to be invented. It needs to be given recoverable structure. The challenge is to impose a traversal path through a large distributed semantic manifold without destroying the generative richness that produced it.

This is not an unusual situation for theoretical work. Many research programs begin as distributed conceptual ecologies before crystallizing into formal objects. What is unusual here is the degree to which the distributed form is itself a consequence of the theoretical commitments. A framework that argues for local reconstruction over hidden abstraction, for provenance over compression, for trajectory visibility over seamless continuity, will tend to produce documents that resist premature closure.

1.2 Local Coordinate Charts over Shared Geometry

Different companion papers foreground different regions of the same underlying geometry. A reader entering through economics encounters admissibility and trajectory restriction through labor markets and platform extraction. A reader entering through cognition encounters sparse heuristics and reconstruction under uncertainty. A reader entering through cosmology encounters entropy gradients, field relaxation, and persistence constraints. A reader entering through linguistics or game design encounters the same operations under different surface vocabularies.

The coordinate-chart metaphor is intended technically rather than merely decoratively. In differential geometry, a manifold may require multiple overlapping coordinate charts to be fully described, with no single chart covering the entire space without singularity. The transition maps between charts specify how descriptions in overlapping regions must agree. The global manifold is not any individual chart but the compatibility structure across all of them.

The conceptual manifold of this monograph has exactly this structure. No single essay or chapter covers the entire space without distortion. The monograph functions as a set of transition maps, establishing how the descriptions in overlapping domains must be reconciled. Where the cosmological and cognitive formulations appear to describe different phenomena, the transition maps reveal that they are coordinate representations of the same underlying invariants.

Definition 1.1 (Semantic Coordinate Chart)

A semantic coordinate chart is a local theoretical framework that provides a consistent description of a region of conceptual space. A chart is *valid* in its region if its internal commitments are non-contradictory and its key terms are operationally grounded. Two charts are *compatible* on their overlap if there exists a translation scheme that preserves inferential relationships across the shared region. A collection of charts constitutes a *semantic atlas* if every region of the conceptual space is covered by at least one valid chart and all overlapping charts are pairwise compatible.

Derivation: How Constraints Generate a Filtered Trajectory Manifold

To make the admissibility concept concrete before it is generalized across domains, consider a toy derivation. Let $\mathcal{X} = \mathbb{R}^n$ be a configuration space and let $\mathcal{C} = \{c_1, \ldots, c_k\}$ be a set of local compatibility conditions, each of the form $c_j(\mathbf{x}, \dot{\mathbf{x}}) \leq 0$. These conditions may be thermodynamic (energy cannot increase above a threshold), logical (contradictory states cannot coexist), geometric (the trajectory cannot pass through excluded regions), or historical (a prior commitment rules out certain futures).

The unconstrained trajectory space $\mathcal{X}^{[0,T]}$ contains all continuous paths $\gamma : [0, T] \to \mathcal{X}$. The admissible subspace is:

$$\mathcal{A}(\mathcal{C}) = \left\{ \gamma \in \mathcal{X}^{[0,T]} : c_j(\gamma(t), \dot{\gamma}(t)) \leq 0 \text{ for all } j \text{ and a.e. } t \in [0, T] \right\}.$$

Three observations follow immediately. First, $\mathcal{A}(\mathcal{C})$ is strictly smaller than $\mathcal{X}^{[0,T]}$ whenever any constraint is non-trivial. Second, the intersection of constraints is itself a manifold (under smoothness conditions), so the admissible space has geometric structure that the unconstrained space does not. Third, adding a constraint always contracts $\mathcal{A}$ and never expands it. Constraint is therefore irreversibly selective.

The key contrast with optimization is that optimization searches $\mathcal{X}^{[0,T]}$ with a cost function and finds a minimum; constrained navigation never considers inadmissible paths at all. The cognitive significance of this distinction is substantial: a system that never enters the inadmissible region need not represent it. Constraint operates before evaluation, not after.

The claim advanced across all companion documents is that the frameworks of RSVP field theory, KES synthesis, Spherepop event calculus, TARTAN tiling, CLIO constraint repair, and the Yarncrawler polycompiler constitute a compatible semantic atlas over a common underlying geometry, even when they appear to address incommensurable subject matters.

1.3 The Compression Problem in Academic Writing

Conventional academic writing assumes stable disciplinary priors and compresses inferential chains aggressively. This produces efficient communication within closed interpretive communities at the cost of accessibility, provenance, and reconstructibility. A single undefined term can silently import an entire conceptual framework. If that framework is not reconstructed locally, readers who do not already inhabit the same intellectual environment are forced to guess the intended meaning.

Many technical communities normalize this compression because it reduces repetition and signals group membership. Entire chains of assumptions become condensed into specialized vocabulary. That increases communication efficiency within a closed interpretive community but decreases accessibility and often obscures the actual inferential structure.

This monograph operates under the opposite constraint. Terms are introduced progressively, derived from previously established concepts, or explicitly flagged as imported from identified external frameworks. Appendices preserve the boot sequence rather than presenting only the compiled executable abstraction. The result is longer and more repetitive than a compressed academic text, but also more legible and less dependent on insider intuition.

> *The analogy to computational bootstrapping is deliberate and will be developed formally in Chapter 12. An operating system does not appear instantaneously from nothing. Firmware initializes primitive hardware assumptions. Bootloaders establish minimal execution environments. Kernels initialize memory structures. Higher-level abstractions only become meaningful because lower-level structures have already been stabilized. The chapter organization of this monograph is designed to enforce the dependency ordering rather than present conclusions before premises.*

1.4 Admissibility as the Primitive Organizing Concept

Before any domain-specific application appears, the notion of admissibility requires introduction. It is the single concept that recurs most persistently across all the frameworks assembled in this monograph.

Definition 1.2 (Admissible Trajectory)

Let $\mathcal{X}$ be a state space and let $\gamma : [0, T] \to \mathcal{X}$ be a continuous path through that space. The path γ is *admissible* with respect to a constraint set $\mathcal{C}$ if it satisfies the local compatibility conditions imposed by $\mathcal{C}$ at every point along its length. Constraints may be thermodynamic, logical, geometric, historical, or any combination thereof. The *admissible trajectory space* $\mathcal{A}(\mathcal{C}) \subseteq \mathcal{X}^{[0,T]}$ is the set of all admissible paths.

Intelligence, computation, and physical evolution are all subsequently treated as processes of navigating admissible trajectories rather than optimizing over unconstrained possibility spaces. The emphasis falls on constraint as productive rather than merely restrictive. Constraints define the space of possibilities worth considering. A system that could traverse any trajectory would have no basis for selection; it is precisely the inadmissibility conditions that make directed behavior possible.

This is why the projects across the corpus repeatedly return to geometry, compression, sparsity, trajectories, irreversibility, and reconstruction. The central intuition is that intelligence does not operate through exhaustive representation of reality. Biological and computational systems survive by constructing partial, compressed, energetically tractable approximations that preserve enough structural invariants to continue operating successfully under uncertainty.

Chapter 2

Core Vocabulary

> Terminology is prevented from becoming detached from operational meaning. The insistence on repeatedly rebuilding assumptions from the ground upward functions as a defense against proxy drift at the conceptual level itself.

This chapter establishes the definitional scaffold on which the remainder of the monograph is built. Terms introduced here are used without redefinition in subsequent chapters. Where a concept is partially anticipated in Chapter 1, the definition here provides the formal completion.

2.1 State Spaces and Trajectory Selection

Definition 2.1 (State Space)

A *state space* $\mathcal{X}$ is a set whose elements represent all possible configurations of a system. A *trajectory* through $\mathcal{X}$ is a map $\gamma : I \to \mathcal{X}$ from an index set I (typically a time interval) into the state space. The *selection pressure* acting on a trajectory space is the set of mechanisms that filter admissible from inadmissible paths. Selection pressure may arise from physical laws, thermodynamic constraints, logical requirements, historical boundary conditions, or any combination of these.

The distinction between exhaustive search and constrained navigation is fundamental. A system performing exhaustive search considers the full space $\mathcal{X}^I$ of possible trajectories and evaluates each against an objective function. A system performing constrained navigation operates only within $\mathcal{A}(\mathcal{C})$ from the outset, never considering inadmissible paths. Biological and cognitive systems almost exclusively employ constrained navigation, because exhaustive search is thermodynamically impossible at any scale of practical interest.

2.2 Compression, Sparsity, and Structural Invariants

Definition 2.2 (Compression)

A *compression* of a system description D is a map $\kappa : D \to D'$ where $|D'| < |D|$ under some measure of description length, and D' preserves a specified set of structural relations present in D. A compression is *lossless with respect to* $\mathcal{R}$ if every relation $r \in \mathcal{R}$ is recoverable from D'. It is *lossy* if some $r \in \mathcal{R}$ is not recoverable. Compression is therefore not reduction of volume alone but selective preservation of relations.

Definition 2.3 (Structural Invariant)

A *structural invariant* of a system under a family of transformations $\{T_\alpha\}$ is a property P such that $P(T_\alpha(s)) = P(s)$ for all α and all states s in the relevant domain. Structural invariants are the properties that compression must preserve to remain operationally useful. A compression that erases structural invariants necessary for understanding, provenance, or accountability is a *destructive compression* regardless of its information-theoretic efficiency.

Definition 2.4 (Sparsity)

A representation is *sparse* if most of its degrees of freedom are zero or near-zero, with significant values concentrated in a small fraction of components. Sparsity is the condition under which compression becomes tractable: a sparse representation can be compressed without destroying structural invariants because most of the space is already empty. The *Natural Sparsity Principle* states that in biological systems, sparsity emerges as a natural consequence of physiological pressures—energetic constraints, signal noise, chemical gradients—without requiring explicit computational penalties to enforce minimal complexity.

Why Sparsity Makes Reconstruction Tractable

The connection between sparsity and tractable reconstruction can be sketched formally. Let $\mathbf{x} \in \mathbb{R}^N$ be a signal with $k \ll N$ nonzero components. The number of bits required to specify $\mathbf{x}$ is approximately $k \log(N/k) + k \log(\max |\mathbf{x}|/\epsilon)$ for precision ϵ, far less than the $N \log(\max |\mathbf{x}|/\epsilon)$ bits required for a dense signal.

Under bounded energetic constraints, a system with budget B (in bits or equivalent energy units) can faithfully represent $\mathbf{x}$ if and only if $B \geq$ the encoding cost. For dense signals this budget is exhausted rapidly; for sparse signals it is not. The compressive sensing literature shows that under sparsity, signals can be recovered from $O(k \log N)$ measurements rather than N measurements, with reconstruction error bounded by the sparsity level.

The biological relevance is direct. Neural firing rates are sparse: typical cortical neurons fire at low average rates, with high-rate bursts carrying most information. This means the energetic cost of neural representation scales with signal sparsity rather than state-space dimensionality. A brain that maintained dense representations of all sensory dimensions simultaneously would require thermodynamically infeasible metabolic support. Sparsity is not a design choice but a thermodynamic necessity that the Natural Sparsity Principle identifies as a universal constraint on cognitive architecture.

2.3 Reconstruction and Hallucination

Definition 2.5 (Reconstruction)

Reconstruction is the process by which a system assembles an operationally coherent configuration from partial traces. A reconstruction is *grounded* if it remains constrained by external feedback signals. It is *corrigible* if errors in reconstruction are detectable and correctable through interaction with the environment. It is *drifting* if it progressively decouples from external constraint and stabilizes on an internally self-consistent but externally inaccurate model.

Definition 2.6 (Hallucination (Formal))

Hallucination is reconstruction that exceeds available constraint: the assembly of a locally coherent configuration from traces that underdetermine it, without external correction signals sufficient to resolve the underdetermination. Hallucination is not a pathology but the default operating condition of any sufficiently compressed system under uncertainty. The critical distinction is not between hallucinating and non-hallucinating systems but between systems whose hallucinations remain *corrigible* through feedback and systems whose internal reconstructions drift irreversibly from external constraint.

> *This definition of hallucination is deliberately broader than its use in AI discourse, where it typically refers to language model confabulation. The formal definition encompasses perception, memory consolidation, theoretical inference, institutional modeling, and any other process in which a partial trace is assembled into a more complete representation. All such processes hallucinate in this sense; the question is always whether the hallucination is correctable.*

Reconstruction Bound

The inevitability of hallucination under compression can be made precise. Let D be a description of cardinality $|D| = N$ and let $\kappa(D) = D'$ be a compression with $|D'| = M < N$. The number of distinct descriptions D consistent with D' is at least 2^{N-M}, since the compressed representation cannot distinguish among them. Any reconstruction process that must select a single completion from among these 2^{N-M} possibilities must hallucinate: it must choose structure that the compression does not uniquely determine.

More formally, if $p(D \mid D')$ is the posterior distribution over original descriptions given the compressed form, then:

$$H(D \mid D') = -\sum_{D} p(D \mid D') \log p(D \mid D') \geq N - M > 0.$$

The conditional entropy is strictly positive whenever $M < N$. A reconstructing system with access only to D' cannot recover D with certainty. It must interpolate the missing $N - M$ bits from prior structure, contextual signals, or statistical regularities. That interpolation is hallucination in the formal sense: reconstruction exceeding available constraint.

The corrigibility condition—the requirement that reconstruction errors remain correctable through feedback—does not eliminate this hallucination. It ensures only that when the interpolated structure turns out to be wrong, the error is detectable and correctable before it propagates into downstream commitments.

2.4 Provenance and the Boot Sequence

Definition 2.7 (Provenance)

The *provenance* of a conceptual or computational object is the recoverable dependency graph of the operations and inputs that produced it. Provenance includes: the sequence of constraint-stabilization events that enabled the object's current form; the assumptions that were imported versus derived; and the points at which external evidence entered the reconstruction process. A system with high provenance transparency is one whose dependency graph can be approximately reconstructed by an external observer.

Definition 2.8 (Boot Sequence)

A *boot sequence* is the ordered set of constraint-stabilization events necessary to make a higher-level abstraction operationally meaningful. Lower-level constraints must be stabilized before higher-level ones can be evaluated. Erasing the boot sequence—presenting only the terminal abstraction without the ordered initialization history—transforms an inspectable object into an opaque one. The boot sequence is thus the provenance of a system's layered structure.

2.5 The Formal Substrates: RSVP, KES, Spherepop

Three formal frameworks function as computational and geometric substrates throughout this monograph. They are introduced briefly here and developed formally in the appendices. References to these frameworks in subsequent chapters should be understood as references to the formal specifications given there.

Definition 2.9 (RSVP Field System)

The *Relativistic Scalar-Vector Plenum* (RSVP) is a coupled field system consisting of three dynamical fields over a spacetime manifold $\mathcal{M}$:

- $\Phi(\mathbf{x}, t)$: a scalar density field encoding concentration, potential, or legitimacy of state.
- $\mathbf{v}(\mathbf{x}, t)$: a vector flow field encoding directed transport, trajectories, and causality.
- $S(\mathbf{x}, t)$: an entropy field encoding disorder, uncertainty, and complexity budgets.

The three fields evolve according to coupled nonlinear partial differential equations. The full Lagrangian structure and seven governing axioms are given in Appendix 13.5. Within the semantic interpretation developed in Chapter 5, Φ measures semantic density, $\mathbf{v}$ encodes recursive trajectories of meaning, and S tracks unresolved noise available for future reinterpretation.

Definition 2.10 (KES Map)

The *Kinetic-Event Synthesis* (KES) map is the function

$$\text{KES} : \Omega_t \longrightarrow H_{t+1}$$

from a current world-state Ω_t to a successor hypothesis space H_{t+1}. The map formalizes the basic cognitive operation: given a current configuration of constraints, evidence, and prior commitments, what hypothesis spaces are admissible at the next step? The full convergence conditions and technical specification are given in Appendix 13.10.

Definition 2.11 (Spherepop Calculus)

Spherepop is an irreversible event calculus defined by four primitive operators acting on a possibility space $\mathcal{P}$:

- Pop: extracts committed events from possibility space, transitioning them to actuality.
- Refuse: eliminates inadmissible trajectories, contracting the admissible space.
- Bind: establishes structural dependencies between events, constraining future possibilities.
- Collapse: commits a possibility space to a definite configuration, rendering alternatives inaccessible.

All four operations are irreversible: once applied, they cannot be undone within the same system. The operational semantics of all four operators are developed formally in Appendix 13.15.

> *The irreversibility of Spherepop operators is not a limitation but a design commitment. Systems that can always reverse their operations accumulate no structural residue; they have no history and therefore no trajectory. The Collapse operator in particular formalizes the commitment structure that distinguishes genuine decision from perpetual deferral.*

Constraint Boundary

What this chapter establishes: Formal definitions of admissibility, compression, sparsity, reconstruction, hallucination, provenance, and the boot sequence. The reconstruction bound $H(D \mid D') \geq N - M > 0$ is formally derived. The three formal substrates (RSVP, KES, Spherepop) are introduced as technical objects.

What this chapter does not establish: That these definitions uniquely capture the relevant phenomena in any particular domain. The definitions are productive insofar as they generate useful inferences; their domain of validity must be established by application.

Remaining speculative: The claim that the Natural Sparsity Principle operates universally across biological systems is a **[C]** conjecture supported by neuroscience and thermodynamics but not formally derived here.

Falsification conditions: A system that produces corrigible reconstruction from a strictly lossless compression would falsify the claim that hallucination is constitutive of compressed inference. A biological cognitive system that maintains dense rather than sparse representations at thermodynamically sustainable cost would falsify the

> Natural Sparsity Principle.

Part II

Cognitive Architecture

> The singularity is less an explosion of raw machine capability and more a transition in representational infrastructure.

Part II develops the Human Cognitive Singularity argument as the central theoretical object of the monograph. Intelligence is treated as a scaffolded amplification process grounded in representational breakthrough rather than metaphysical discontinuity. The Ellul constraint on technical opacity is formalized in the vocabulary of admissibility and trajectory visibility.

Chapter 3

Intelligence as Scaffolded Amplification

> Intelligence advances not merely through raw computational power, but through new methods of organizing, compressing, navigating, and stabilizing complexity.

3.1 Against the Replacement Narrative

The dominant cultural framing of artificial intelligence treats it as a future autonomous species superseding biological intelligence. This framing misidentifies the phenomenon. Intelligence has always advanced through new methods of organizing, compressing, navigating, and stabilizing complexity—not through the emergence of a qualitatively alien substrate that displaces prior cognitive forms.

The replacement narrative implicitly assumes a threshold model: intelligence exists in discrete levels, and crossing a threshold produces a categorically different kind of system. This assumption is empirically unsupported. The history of cognitive development, both biological and cultural, shows continuous scaffolding rather than threshold crossing. New representational capabilities extend the reach of prior ones without rendering them obsolete.

3.2 Evolutionary Continuity and Representational Breakthrough

The history of intelligence reconstructs as a sequence of representational breakthroughs rather than metaphysical discontinuities. Each transition extends the range of admissible trajectories a cognitive system can navigate while preserving enough structural continuity to remain operationally coupled to its environment.

From bacterial chemotaxis—directional movement along chemical gradients—through multicellular coordination, symbolic language, writing, mathematics, and machine learning, the pattern is consistent. Each breakthrough provides a new compression scheme for environmental regularities, a new class of structural invariants that can be tracked across transformations, a new class of admissible trajectories that can be navigated without exhaustive search.

> *Proposition 3.1* (*Representational Breakthrough Condition*)
>
> A transition in representational capability constitutes a cognitive breakthrough if and only if it (i) strictly extends the admissible trajectory space available to a cognitive system—enabling navigation of problem spaces that were previously inaccessible—while (ii) preserving sufficient structural coupling with prior representations that the new capability can be acquired through scaffolded extension rather than wholesale replacement.

The significance of this proposition is that it rules out both the replacement narrative (new systems render prior ones obsolete) and the incremental narrative (all change is merely quantitative). Genuine breakthroughs are qualitative in their effect on the accessible trajectory space but continuous in their developmental pathway.

3.3 Compression as Constitutive Operation

No organism or computational system can represent the full state of reality. Compression is therefore not a limitation imposed on intelligence from outside but the constitutive operation through which intelligence exists at all. Representation always involves selecting which relations to preserve and which to discard.

Definition 3.1 (Constitutive Compression)

Compression is *constitutive* of a cognitive process if the process could not exist without it: if complete, uncompressed representation of the relevant domain would exceed the energetic, temporal, or informational budget available to the system. For any biological or computational system operating under physical constraints, compression of the world-model is constitutive rather than optional.

The danger arises not from compression itself but from compression that erases relations necessary for understanding, provenance, or accountability. A compressed representation that preserves the structural invariants necessary for the system's operational goals is adequate. A compressed representation that erases those invariants while remaining locally coherent produces the proxy stabilization failure mode described in Chapter 4.

3.4 Cognitive Geometry and Spatial Navigability

Human reasoning operates more effectively when structural relationships become spatially navigable rather than remaining hidden inside opaque symbolic chains. Spatial nesting, dependency tracking, trajectory visualization, and visible constraint relations reduce cognitive load because biological intelligence evolved in geometrically structured environments.

This is why the recurring emphasis across companion documents on geometric interfaces, visible process, and navigable semantic space is not aesthetic preference but cognitive ergonomics. Computation and semantic organization are reinterpreted as forms of navigation through constrained possibility spaces rather than symbolic manipulation of static representations.

Definition 3.2 (Cognitive Geometry)

The *cognitive geometry* of a representational system is the spatial or quasi-spatial structure imposed on its content that makes structural relationships navigable rather than merely searchable. A representational system has high cognitive geometry if the distance between related concepts corresponds to cognitive access cost, if structural dependencies are visible as positional relationships, and if the traversal of the representation space matches the natural grain of the problem structure.

3.5 The Singularity as Representational Infrastructure Transition

The human cognitive singularity is not an explosion of raw machine capability. It is a transition in representational infrastructure: the progressive externalization of human cognitive reach into semantic systems capable of preserving structural relationships across scales of complexity that unaided biological cognition cannot directly navigate.

> *Proposition 3.2* (*Amplification Rather Than Substitution*)
>
> The productive framing of artificial intelligence as cognitive infrastructure treats AI systems as prosthetics for human reasoning rather than replacements for it. A prosthetic extends reach while preserving the agent's participation in the process. A replacement eliminates the agent from the loop. The design objective for cognitive infrastructure is therefore maximal amplification of human reasoning capacity with minimal reduction of human participation in the reasoning process itself.

The emphasis falls on amplification rather than substitution. The representational system functions as a prosthetic for human reasoning rather than a replacement for it. The danger is not that AI systems will become too capable but that they will become too opaque: optimization engines that sever humans from the causal and semantic structures they depend on. The hope is that transparent geometric and constraint-oriented representations can instead extend collective reasoning while preserving intelligibility, accountability, and participation.

3.6 Inferential Continuity versus Symbolic Continuity

One of the most consequential distinctions for evaluating representational systems is the difference between *inferential continuity* and *symbolic continuity*. The distinction operationalizes much of what the preceding sections argue about compression and provenance.

Definition 3.3 (Symbolic Continuity)

A representation exhibits *symbolic continuity* if it preserves the surface flow of recognizable language or notation: if each statement is locally coherent, terminologically consistent, and grammatically well-formed. Symbolic continuity is a necessary condition for communicability but is insufficient for epistemic transparency.

Definition 3.4 (Inferential Continuity)

A representation exhibits *inferential continuity* if each statement is derivationally connected to prior statements through explicit steps that can be reconstructed by the reader. Inferential continuity requires that the dependency graph of the argument be recoverable, not merely that the argument be locally plausible.

Symbolic continuity without inferential continuity is the formal definition of what might be called *synthetic fluency*: the production of semantically coherent text that does not preserve derivational constraint across its length. A system optimized for symbolic continuity will produce outputs that appear rigorous at the sentence level while accumulating inferential drift at the paragraph and chapter level.

This is empirically observable in large language models. Systems trained to predict the next token under a plausibility objective will develop high symbolic continuity—they produce text that reads like academic prose—while exhibiting low inferential continuity: the argument of paragraph five may be incompatible with the argument of paragraph two without any local signal that a contradiction has occurred. Each sentence is locally admissible; the global trajectory is not.

The same failure mode appears in institutional discourse, bureaucratic documentation, and theoretical writing that has been compressed across many rewrites. The surface vocabulary remains consistent while the inferential structure quietly degrades.

[P] Provenance transparency, as defined in Chapter 2, is the mechanism by which inferential continuity is maintained under compression. A representation with high provenance transparency allows the reader to trace each claim back through its dependency sequence to the primitive constraints that ground it. The boot sequence is inspectable. A representation with low provenance transparency presents only the terminal abstraction, leaving the derivational structure hidden.

The monograph's own style—repetitive definition accumulation, explicit cross-referencing, refusal to assume disciplinary priors—is a direct consequence of prioritizing inferential over symbolic continuity. It produces a text that is slower and more redundant than compressed academic writing, but whose dependency structure remains visible throughout.

> **What this chapter establishes:** The representational breakthrough account of intelligence as scaffolded amplification; the formal definitions of symbolic and inferential continuity; the prosthetic framing of AI as amplification rather than replacement.
>
> **Remaining conjectural [C]:** The claim that all cog-

nitive breakthroughs follow the representational breakthrough pattern is a conjecture. The pattern is well-supported across documented transitions (language, writing, mathematics) but the sample is small and the selection criteria are unclear.

Falsification conditions: A cognitive breakthrough that dramatically extends accessible trajectory space without scaffolded extension from prior representations would falsify the continuity thesis. The emergence of written language ex nihilo with no prior symbolic system would be such a case.

Chapter 4

The Ellul Constraint

> Once compressed representations become detached from the systems they originally tracked, optimization begins operating against the proxy rather than the underlying reality. Metrics become targets. Simulations become substitutes.

4.1 Technical Systems and the Opacity Problem

Jacques Ellul's analysis of technique tepellul1964 identifies a structural tendency in modern technical infrastructures toward opacity, massification, and the conversion of individuals into components of large-scale optimizing processes. Technical systems privilege efficiency and predictability over interpretability and human agency. The Ellul critique is not merely sociological; it identifies a structural failure mode that can be formalized in the vocabulary of admissibility and trajectory visibility.

The failure mode is proxy stabilization: a compressed representation of a system becomes operationally substituted for the system itself, and optimization processes subsequently operate on the proxy rather than the underlying reality.

Definition 4.1 (Proxy Stabilization)

Proxy stabilization occurs when a compressed representation R of a system $\mathcal{S}$ becomes decoupled from $\mathcal{S}$ through the following sequence: (i) R is constructed as a useful compression of $\mathcal{S}$ at time t_0; (ii) optimization processes are applied to R rather than $\mathcal{S}$ directly, because $\mathcal{S}$ is inaccessible or too expensive to evaluate; (iii) the optimization pressure on R changes R in ways that diverge from changes in $\mathcal{S}$; and (iv) R stabilizes in a configuration that is internally self-consistent but externally inaccurate. The result is a drifting reconstruction in the sense of Definition **??**.

4.2 Proxy Stabilization and Metric Drift

Proxy stabilization produces metric drift: the metric that was intended to track a property of the underlying system instead tracks a property of the proxy, which has been optimized to score well on the metric regardless of the underlying property. This is Goodhart's Law stated in terms of the trajectory framework:

when a measure becomes a target, it ceases to be a good measure, because optimization pressure drives the system toward the proxy attractor rather than the underlying attractor.

The same failure mode appears across radically different domains. Academic citation counts are proxies for research quality that become targets when hiring decisions optimize for them. Engagement metrics are proxies for user satisfaction that become targets when recommendation algorithms optimize for them. Body mass index is a proxy for health that becomes a target when insurance pricing optimizes for it. In each case, the proxy initially tracks something real, then becomes decoupled from it through optimization pressure, then substitutes for it in downstream processes.

> *Proposition 4.1* (*Proxy Drift Theorem*)
>
> Let $\mathcal{S}$ be a target system, R a compressed representation of $\mathcal{S}$, and $\mathcal{O}$ an optimization process operating on R. If (i) $\mathcal{O}$ has sufficient capacity to find configurations of R that score well on its objective independently of the state of $\mathcal{S}$, and (ii) the feedback loop from $\mathcal{S}$ to R is slower than the optimization loop of $\mathcal{O}$, then R will drift away from accurate representation of $\mathcal{S}$ at a rate proportional to the optimization pressure and inversely proportional to the feedback bandwidth.

Dynamical Sketch of Proxy Drift

Let $s(t)$ denote the true state of $\mathcal{S}$ and $r(t)$ the proxy representation, with initial agreement $r(0) \approx s(0)$. The optimization process applies gradient updates to r:

$$\dot{r} = -\nabla_r \mathcal{L}(r) + \alpha\,(s - r),$$

where $\mathcal{L}(r)$ is the optimization objective defined over the proxy and $\alpha > 0$ is the feedback bandwidth (the rate at which the proxy is corrected toward the true state). The first term drives r toward minima of $\mathcal{L}$; the second term pulls r toward s.

If $\mathcal{L}$ has a minimum at $r^* \neq s$—which is generically the case once optimization pressure is sufficient to exploit the proxy—then the equilibrium of this system is not $r = s$ but:

$$r^* = s - \frac{1}{\alpha}\,\nabla_r \mathcal{L}(r^*).$$

The drift $\|r^* - s\|$ grows as α decreases (weaker feedback) and as $\|\nabla_r \mathcal{L}(r^*)\|$ increases (stronger optimization pressure). In the limit $\alpha \to 0$ (no feedback from ground truth), r tracks $\mathcal{L}$ alone and decouples entirely from s.

This is not a failure of the optimization process. The optimization process is working correctly by its own lights: it is minimizing $\mathcal{L}$ over r. The failure is structural: the proxy r was assumed to track s, but the assumption breaks down under optimization pressure faster than feedback can correct it.

4.3 Corrigibility, Feedback, and Constraint Recovery

The response to proxy drift is not rejection of compression but the design of systems whose internal reconstructions remain constrained by external feedback rather than self-referentially stabilizing.

Definition 4.2 (Corrigibility)

A system is *corrigible* with respect to a target domain if its internal model of that domain remains responsive to constraint signals from the domain itself. Corrigibility is a structural property: it requires that the feedback loop from external reality to internal representation remain open and operative, with sufficient bandwidth to correct drift before it becomes self-reinforcing. A system that has achieved proxy stabilization has lost corrigibility with respect to the proxied domain.

Corrigibility distinguishes adaptive sparse heuristics from pathological proxy optimization. Sparse heuristics are compressed representations that remain corrigible: they simplify the world-model aggressively while preserving the feedback mechanisms that detect and correct reconstruction errors. Pathological proxies achieve compression by severing those feedback mechanisms.

4.4 Cognitive Geometry as the Design Response

The design alternative to opaque statistical optimization is systems organized around visible structure, provenance, interpretable dependencies, and navigable semantic relations. This is the positive proposal that runs across the educational tools, audio visualizers, geometric interfaces, and symbolic scripts developed in companion work: representational infrastructure designed to extend human reasoning while preserving the intelligibility, accountability, and participation that opacity forecloses.

The recurring CRT aesthetics, oscilloscope displays, phosphor trails, and recursive traces in the companion software are not decorative. They externalize process. Instead of hiding computation behind polished interfaces, they expose trajectories, residue, decay, stabilization, weighting, reconstruction, and progressive constraint selection as visible operations. The interface is designed to make the system's internal dynamics inspectable rather than seamless.

This connects directly to the provenance requirement. A system optimized solely for final answers loses the structure necessary for understanding how those answers emerged. A system designed around visible trajectories preserves that structure, allowing the reconstruction process to be approximately replayed by an observer even when the original process has terminated.

Constraint Boundary

> **What this chapter establishes:** The proxy stabilization failure mode as a structural dynamic, formalized via the differential equation $\dot{r} = -\nabla_r \mathcal{L}(r) + \alpha(s - r)$. The Proxy Drift Theorem shows that drift magnitude scales with optimization pressure and inversely with feedback bandwidth. Corrigibility is defined formally.
>
> **What this chapter does not establish:** That all proxy drift is malicious or even harmful. Some proxy stabilization produces useful compression without meaningful decoupling. The theorem describes a failure mode, not an inevitable outcome.
>
> **Remaining phenomenological [P]:** The applications to academic citation counts, engagement metrics, and institutional measurement are phenomenological interpretations of the formal result, not formal derivations from it.
>
> **Falsification conditions:** A compressed representation that remains stably coupled to its target system under

> sustained optimization pressure, without requiring bandwidth α greater than the optimization rate, would falsify the theorem's claim that drift is structurally inevitable under those conditions.

Part III

Yarncrawler: Self-Repair as Universal Principle

> To act intelligently is to Yarncrawl: to repair oneself by weaving meaning into structure, to sustain homeorhetic flows rather than static equilibria, and to treat skepticism not as a threat but as a guide to resilience.

Part III develops the Yarncrawler framework as the self-repair layer of the entire theoretical architecture. The Yarncrawler is not an application of the foregoing concepts but the mechanism by which they maintain coherence with each other. The sheaf-theoretic foundation of Chapter 6 provides the formal statement of what the monograph is doing to itself: stitching overlapping coordinate charts into a global section without destroying the generative richness of the local descriptions.

Chapter 5

Yarncrawler: The Self-Refactoring Polycompiler

> A Yarncrawler does not simply interpret its world—it keeps itself alive by reinterpreting itself into being.

5.1 The Operative Metaphor and Its Formal Content

The Yarncrawler framework models organisms, cultural systems, and artificial intelligences as stigmergic parsers maintaining homeorhetic Markov blankets. The name is not merely evocative. A crawler moves through an environment by traversal. A yarn system is structured through threads that can be unwound and rewound. A polycompiler processes multiple input formats and compiles them into a unified output while modifying its own compilation rules in response to what it encounters. The combination captures the essential operation: a system that parses its environment while simultaneously rewriting its own grammar.

A useful image is that of a ball of yarn with a guiding aperture, a semantic straw. Threads are continuously pulled from the interior, passed through the aperture, reinforced with contextual material, and rewound on the exterior. This cycle of extraction, transformation, and reintegration parallels how agents parse inputs, re-thread outputs, and repair their own structural coherence.

5.2 The Seven Swarm-Care Axioms

Before the formal definition, the operational principles that govern Yarncrawler behavior are stated as axioms. These axioms were developed in the framework's original formulation and provide the semantic content that the mathematical definition subsequently formalizes.

Axiom 5.1 (Local Action Only)

A Yarncrawler acts only on its current node and its local neighbors. No global inference is required or permitted. Under the Markov blanket decomposition, local conditional independence guides node-specific belief revision: only local evidence affects local updates.

Axiom 5.2 (Semantic Trail Documentation)

All operations must be logged: input state, evaluation, action taken, and outcome. Every correction leaves a trace—a map of belief revision tied to the confidence in the prediction that was corrected.

Axiom 5.3 (Autocatalytic Fix Design)

Fixes should ease future inference and adaptation. A repair operation is *autocatalytic* if it lowers the cost of subsequent repairs in the same region. Good fixes are not merely locally optimal—they are self-amplifying priors that reshape the generative model.

Axiom 5.4 (Trail Reinforcement and Decay)

Successful trails are strengthened; failing trails are allowed to fade. Trail weights evolve according to $w_{t+1} = w_t \cdot \exp(-\lambda \Delta t)$ for decay and $w_{t+1} = w_t + \alpha R(x_t)$ for reinforcement, where R is a success signal and α a learning rate.

Axiom 5.5 (Node Hygiene)

A Yarncrawler maintains its environment even when conditions appear stable. Preventative maintenance is inference under uncertainty: the system anticipates disorder and reduces entropy before it accumulates past the repair threshold.

Axiom 5.6 (Failure Isolation and Recovery)

When a node exceeds a critical damage threshold, it is quarantined. Its state snapshot is archived for future re-entry. Retreat and data capture are rational under epistemic overload. The system does not attempt repair beyond its current capacity.

Axiom 5.7 (Swarm Feedback Integration)

Yarncrawlers share belief updates across the swarm. The combined belief is formed by precision-weighted averaging: $q_{\text{combined}}(s) = \sum_i \Pi_i q_i(s) / \sum_i \Pi_i$, where Π_i is the precision (inverse uncertainty) of crawler i's update. Confidence acts as epistemic gravity.

5.3 Formal Definition

Definition 5.1 (Yarncrawler)

A *Yarncrawler* is a quintuple

$$\mathcal{Y} = \big(\mathcal{M},\ \{(U_i, \varphi_i, f_i)\}_{i\in I},\ \{w_i\}_{i\in I},\ \mathcal{K},\ \mathcal{R}\big)$$

where:

- $\mathcal{M}$ is a smooth semantic manifold with atlas $\{U_i\}_{i\in I}$;
- $\varphi_i : U_i \to \mathbb{R}$ is the local scalar potential (semantic density) of expert i;
- $f_i : U_i \to T\mathcal{M}$ is the local vector field (directed repair flow) of expert i;
- $\{w_i\}_{i\in I}$ is a partition of unity subordinate to $\{U_i\}$, so that $\sum_{i\in I} w_i(x) = 1$ and $\operatorname{supp}(w_i) \subseteq U_i$;
- $\mathcal{K}$ is a finite knowledge store with embedding φ_{emb};
- $\mathcal{R}$ is a stigmergic reinforcement operator updating the gates w and local fields (φ, f) subject to minimization of a seam penalty and free-energy functional.

The global RSVP fields are defined from these local data by:

$$\Phi(x) = \sum_i w_i(x)\,\varphi_i(x), \qquad \mathbf{v}(x) = -\nabla\Phi(x), \qquad S(x) = -\sum_i w_i(x) \log$$

The dynamics of a Yarncrawler trajectory x_t are:

$$\dot{x}_t \in \mathrm{co}\{f_i(x_t) : w_i(x_t) > 0\} - \nabla\Phi(x_t) + B(x_t)\,u_t,$$

where B encodes blanket-mediated control and co denotes the convex hull.

Worked Interpretation: What Each Component Means

The mathematical objects in the quintuple map onto concrete behaviors as follows.

$\mathcal{M}$ is the space of possible semantic states the system can occupy: all the configurations of knowledge, belief, and parsed structure that are available to it. For a squirrel, $\mathcal{M}$ is the space of foraging strategies indexed by environmental cues. For a language model, it is the latent space of its semantic representations.

The local expert pairs (U_i, φ_i, f_i) are the system's specialized competence regions. A squirrel has different behavioral routines for open meadows, dense undergrowth, and known cache sites; each routine is a local expert valid in its region U_i. The scalar potential φ_i measures how strongly the expert "claims" a given state as its domain of competence. The vector field f_i specifies which direction to move within that domain.

The partition of unity $\{w_i\}$ is the gating mechanism. At any given state, multiple experts may have partial competence. The weights $w_i(x)$ determine how much each expert contributes to the current behavior. The entropy $S(x) = -\sum_i w_i \log w_i$ measures how spread the competence is: low entropy means one expert dominates; high entropy means multiple experts are equally relevant, which corresponds to genuine semantic ambiguity.

The knowledge store $\mathcal{K}$ is the stigmergic memory: the accumulated traces of prior traversals. In the squirrel, it is the system of cached seeds and the spatial memory of where they are. In the Bruniquel cave builders, it was the accumulating knowledge of which stalagmite lengths produce which frequencies.

The reinforcement operator $\mathcal{R}$ is the repair mechanism. After each traversal, it adjusts w_i, φ_i, and f_i to reduce the seam penalty: the mismatch between adjacent experts at their boundaries. This is the self-refactoring operation that gives the Yarncrawler its name. The grammar of the parser is rewritten by the act of parsing.

5.4 The Yarncrawler as a Lax Monoidal Functor

The category-theoretic formulation provides the cleanest single-sentence characterization of the Yarncrawler.

Definition 5.2 (Yarncrawler Parser (Category-Theoretic))

Let $\mathcal{W}$ be the *world-trajectory category* whose objects are finite sensorimotor traces $\tau_{[t_0,t_1]}$ and whose morphisms are temporal concatenations and coarse-grainings. Let $\mathcal{M}$ be the *internal semantic category* whose objects are typed semantic modules and whose morphisms are admissible rewrites. Let **Reconf**$(\mathcal{M})$ be the reconfiguration 2-category encoding self-repair moves as 2-morphisms.

A *Yarncrawler parser* is a lax monoidal functor

$$\mathsf{P} : (\mathcal{W}, \otimes) \longrightarrow (\mathcal{M}, \boxtimes)$$

together with a self-rewrite endofunctor $\mathsf{R} : \mathcal{M} \to \mathcal{M}$ such that for each trace τ,

$$\mathsf{P}(\tau) \xRightarrow{\eta_\tau} \mathsf{R}(\mathsf{P}(\tau))$$

is a 2-cell in **Reconf**$(\mathcal{M})$ satisfying locality (only modules near the current "tear" are modified) and type safety (repairs remain within the admissible semantic category).

> *The lax rather than strict monoidal structure is essential. A strict functor would require that the parsing of concatenated traces be exactly the concatenation of their parsings. Laxness permits the coherence maps to be non-trivial: the parsing of a sequence can differ from the sequential composition of parsings, which is precisely what happens when the self-repair operation* R *modifies the parser midway through a trace.*

5.5 Homeorhetic Viability

Definition 5.3 (Homeorhetic Viability Band)

A closed set $\mathcal{V} \subseteq \mathcal{M}$ is *viable* for a Yarncrawler if for every initial state $x_0 \in \mathcal{V}$ there exists a control u_t such that $x_t \in \mathcal{V}$ for all t and

$$\int_0^\infty \big(\beta\, L_{\text{seam}}(x_t) + \sigma(x_t)\big)\, dt < \infty,$$

where L_{seam} is the seam penalty measuring mismatch between adjacent local experts and $\sigma(x_t)$ is the entropy export rate. The system is *homeorhetically stable* if $\mathcal{V}$ is viable and the seam penalty converges to zero along optimal trajectories.

Homeorhesis—the maintenance of a stable trajectory rather than a stable equilibrium—is the appropriate stability concept for Yarncrawler systems. Such systems do not seek rest states; they seek the continuation of viable process. The distinction matters for the cosmological projection of Chapter 8: the universe, on the RSVP interpretation, is homeorhetically stable in its entropy relaxation trajectory rather than asymptotically approaching a fixed equilibrium state.

Chapter 6

Sheaf-Theoretic Foundations of Semantic Repair

> The Yarncrawler is a sheaf-based repair machine. Local parsing windows correspond to stalks, repair corresponds to introducing new morphisms, and entropy corresponds to cohomological obstruction.

6.1 Semantic Space and Covers

Let X denote a semantic space—the domain of cultural landmarks, neural states, semantic nodes, or any other structured collection of meaning-bearing objects. A cover $\mathcal{U} = \{U_i\}_{i \in I}$ represents overlapping neighborhoods of X, each corresponding to a local parsing window within which semantic content can be coherently described.

Toy Example: Three Overlapping Interpretive Windows

Before the general definition, consider the simplest case that exhibits a gluing failure. Let $X = [0, 3]$ represent a temporal sequence of observations, covered by three overlapping windows:

$$U_1 = [0, 1.5], \quad U_2 = [1, 2.5], \quad U_3 = [2, 3].$$

Each window U_i contains a local section s_i: a semantic interpretation of what was observed during that interval. Suppose:

- s_1 concludes: the object encountered is a seed.
- s_2 concludes: the object encountered is a pebble.
- s_3 concludes: the object is food.

On the overlap $U_1 \cap U_2 = [1, 1.5]$, both s_1 and s_2 must agree about what was happening during that interval. But "seed" and "pebble" are incompatible. This incompatibility is a *Čech 1-cocycle*: a mismatch on an overlap that prevents gluing into a consistent global interpretation.

A repair operation ρ_{12} introduces the softened category "seed-like object" into s_2, which is compatible with both. The cocycle is trivialized: the 1-cocycle becomes a coboundary. On the overlap $U_2 \cap U_3 = [2, 2.5]$, "seed-like object" and "food" are compatible under the coarser category "edible item." A global section s on all of $[0, 3]$ now exists: the observation was of an edible, seed-like item.

Crucially, the system defers collapse. Until U_3 is observed, multiple global sections remain viable ("pebble" vs. "seed-like"). The non-uniqueness of the global section is captured by the non-triviality of H^1. Only when U_3 provides the additional constraint "food" does the ambiguity resolve. Strategic ambiguity—holding multiple viable global sections in suspension—is the correct epistemic response to underdetermined gluing.

Definition 6.1 (Presheaf of Semantic Modules)

A *presheaf of semantic modules* is a contravariant functor

$$\mathcal{S} : \mathcal{U}^{\mathrm{op}} \to \mathbf{Cat}$$

that assigns to each open neighborhood $U \subseteq X$ a category $\mathcal{S}(U)$ of local semantic modules. Objects of $\mathcal{S}(U)$ are local semantic threads, claims, cached structures, or repaired routines valid within U. Morphisms are semantic restrictions or rewrites. For inclusions $V \subseteq U$, the restriction functor $\rho^U_V : \mathcal{S}(U) \to \mathcal{S}(V)$ encodes context reduction: what is asserted in the larger context may be forgotten or specialized in the smaller one.

6.2 The Sheaf Condition and Semantic Repair

Definition 6.2 (Semantic Tear and Repair)

A family of local sections $\{s_i \in \mathcal{S}(U_i)\}$ is *glueable* if: (i) local consistency holds: $\rho^{U_i}_{U_i \cap U_j}(s_i) \cong \rho^{U_j}_{U_i \cap U_j}(s_j)$ for all overlapping pairs; and (ii) there exists a global extension $s \in \mathcal{S}(\bigcup_i U_i)$ restricting to each s_i.

A *semantic tear* occurs when local consistency fails on some overlap $U_i \cap U_j$: the restriction of s_i and the restriction of s_j to their common domain are not isomorphic. A *repair* is the introduction of new objects or morphisms into $\mathcal{S}(U_i)$ or $\mathcal{S}(U_j)$ such that glueability is restored.

6.3 Cohomology as Semantic Entropy

Cohomology detects the unresolved incompatibilities that remain after local repairs have been applied. The Čech cohomology groups of the presheaf $\mathcal{S}$ with respect to the cover $\mathcal{U}$ organize these obstructions by dimension:

$$\begin{array}{rl} H^0(X, \mathcal{S}) & \text{viable global sections (coherent meanings)}, \\ H^1(X, \mathcal{S}) & \text{minimal ambiguities (semantic seams)}, \\ H^k(X, \mathcal{S}),\ k \geq 2 & \text{deeper obstructions to coherence.} \end{array}$$

These cohomology groups are interpreted as measures of semantic entropy: the unresolved cocycles represent ambiguity stored for potential reinterpretation. Repair corresponds to introducing new morphisms that trivialize cocycles, moving from H^1 toward H^0 by resolving the incompatibilities that the cocycles record.

> *Strategic ambiguity is a resource rather than a failure. A Yarncrawler that collapses all ambiguity immediately—forcing a unique global section at every step—forfeits the adaptability that deferred gluing provides. The presence of a non-trivial H^1 is a signal that multiple viable interpretations exist; the system defers collapse until additional evidence arrives to determine which interpretation to commit to.*

6.4 The Stigmergic Repair Closure Theorem

Theorem 6.1 (Stigmergic Repair Closure)

Let $\mathcal{S}$ be a presheaf of local semantic sections over patches $\{U_i\}$, with overlaps $U_i \cap U_j$ that may admit inconsistent interpretations. Suppose that for each overlap there exists an elastic repair operation ρ_{ij} satisfying:

1. ρ_{ij} preserves type safety: objects remain within a coarser semantic category (for example, "seed-like" rather than either "seed" or "pebble");
2. ρ_{ij} reduces local free energy $\mathcal{F}(U_i \cap U_j)$ relative to the unrepaired sections; and
3. ρ_{ij} is stigmergic: the cost of subsequent repairs decreases with the density of prior contributions.

Then the family of repaired sections $\{s_i^\rho\}$ admits at least one global section s with $\mathcal{F}(s) \leq \sum_i \mathcal{F}(s_i)$ and with viability increasing in the cumulative density of repairs.

> *Proposition 6.1* (*Repair Preserves Blanket Structure*)
>
> If every repair is generated by local morphisms that commute with restrictions and respect conditional independence at the Markov blanket—that is, if $X_t \perp\!\!\!\perp E_t \mid (S_t, A_t)$ is maintained throughout—then any finite composition of repairs preserves the gluing structure of $\mathcal{S}$.

6.5 Self-Maintaining Growth under Stigmergic Feedback

Theorem 6.2 (Self-Maintaining Growth)

Let $\Phi(t)$ denote the accumulated density of a repaired structure (cache, berm, semantic section). Suppose the repair cost function $C(n)$ for the n-th contribution is monotonically decreasing in cumulative density $\Phi(n)$. Then there exists a critical threshold Φ^* such that

$$\Phi(n+1) - \Phi(n) \geq \delta > 0 \quad \text{whenever } \Phi(n) \geq \Phi^*,$$

that is, once Φ^* is surpassed, the system enters a regime of self-sustaining growth where each new repair lowers the barrier for future repairs.

The discrete-time dynamics implementing this theorem take the form:

$$\Phi_{t+1} = \Phi_t + \Delta t \left(\alpha \frac{\Phi_t}{\Phi_t + \theta} \left(1 - \frac{\Phi_t}{K} \right) - \delta \Phi_t \right),$$

where $\alpha \equiv \lambda\kappa$ is baseline inflow, $\theta > 0$ encodes early friction (the cost of initial contributions), K is carrying capacity, and $\delta \geq 0$ is decay. The stigmergic threshold condition is $\alpha > \delta\theta$.

6.6 Ecological and Cultural Instantiations

The abstract sheaf-repair mechanism is grounded by two concrete instantiations that run across the ecological and cultural domains analyzed in the companion work.

In squirrel foraging, each foraging episode is a local patch U_i with a section s_i encoding the interpretation of an encountered object. A nut recognized in one context and a visually similar pebble encountered in another produce conflicting restrictions at the overlap. Rather than collapsing, the squirrel repairs by adopting a softened category—"seed-like object." The cocycle condition fails strictly but is restored under a relaxed gluing criterion, preserving viability. Multiple possible global interpretations are held in suspension until additional evidence resolves the ambiguity.

In berm construction, an analogous process occurs at the cultural scale. A villager encountering a cairn at a territorial edge may perceive it as insufficient and add more material. The material composition need not be uniform; bones, ash, and dung may be included alongside rocks. Local repairs introduce semantic elasticity: the criterion is “boundary-marker-like” rather than “stone-only.” Over time, overlapping deposits cohere into self-maintaining structures that attract further contributions, forming the Voronoi-like tessellations of Amazonian terra preta.

In both cases, the global section—the food cache or the territorial berm—emerges not from top-down planning but from the gluing of locally inconsistent but compatible acts of repair. This is the universal mechanism of the Yarncrawler framework: local disagreements reconciled through semantic elasticity, strategic ambiguity, and stigmergic reinforcement.

Chapter 7

Ecological, Cultural, and Computational Projections

> Across ecological, cultural, and artificial domains, the invariant dynamic is stigmergic reinforcement coupled to Markov blanket repair.

7.1 The Spectrum from Squirrel to Language Model

The Yarncrawler principle is that agents do not merely consume; they weave, cache, and repair. Each action leaves a trace that both satisfies immediate needs and sets the stage for recursive self-maintenance. This dynamic recurs across ecological, cultural, and computational scales. The invariants under this scale change are the RSVP field correspondences: Φ measures accumulated deposits, $\mathbf{v}$ encodes recursive flows, and S quantifies robustness.

7.2 Ecological Yarncrawlers

A squirrel does not carry a metric map of its forest. Its survival depends on weighted affordances: gradients of safety, concealment, and food density blurred across its perceptual field. When hungry, it follows Gaussian-blurred gradients toward likely caches rather than performing global optimization. Its over-provisioning of seeds is not an error but the creation of semantic attractors—distributed caches that stabilize its future viability.

The perceived affordance field is $\psi(x) = (G_\sigma * u)(x)$ where $u(x)$ is the distribution of hidden resources and G_σ is a Gaussian smoothing kernel. Foraging dynamics are $\dot{x}(t) = -\nabla\psi(x) + \eta(t)$. Each cache deposit increments $u(x_k)$, making the field the squirrel uses simultaneously a navigation tool and a stigmergic memory. The forest itself becomes a distributed memory field, with excess ensuring robustness against local failure.

7.3 Cultural Yarncrawlers

Cultural Yarncrawlers instantiate the same principles at a larger scale. In Amazonian communities, household waste was exported to boundary piles that grew fertile through microbial and chemical processes. Once above a critical threshold, these berms self-maintained, attracting further deposition and forming Voronoi-like tessellations of territory. The resulting terra preta soils persist for millennia.

The governing dynamics are analyzed in the companion paper *Yarncrawler in Action* through a system of reaction-advection-diffusion equations for berm biomass B, char matrix C, and nutrients N. The stigmergic threshold condition $\alpha_B \langle S \rangle_{A_\tau} > \delta_B \theta_{\text{eff}}$ formalizes the transition from fragile deposition to self-maintaining growth.

7.4 Artificial Yarncrawlers

Mixture-of-Experts architectures and retrieval-augmented generation instantiate the Yarncrawler principle computationally. Each expert module functions as a local chart on the semantic manifold. Retrieval mechanisms act as stigmergic glue: a memory trace that is frequently retrieved becomes more salient, just as a cairn or berm attracts additional reinforcement. The global semantic field is assembled from local expert densities by $\Phi(x) = \sum_i w_i(x) \varphi_i(x)$.

The key distinction from standard MoE/RAG architectures is that the repair mechanisms of a Yarncrawler reconfigure the modules themselves rather than only gating access to them. The semantic structure is continuous and trajectory-based, modeled by sheaf gluing and RSVP fields, rather than discrete or static.

7.5 The Climate Response as Worked Example

The read-evaluate-write-move loop provides the clearest operational instantiation of the KES map $\Omega_t \to H_{t+1}$ anywhere in the corpus. Consider a Yarncrawler deployed across a coastal region:

Node 1 (Seattle): Reads rising sea levels and failing levees. Evaluates that current barriers are underfunded. Writes a plan to redirect local taxes to infrastructure. Moves to Node 2, carrying the tax-redirection strategy as a candidate fix for other urban nodes.

Node 2 (Olympic National Forest): Reads increased wildfire risk. Evaluates that invasive species are worsening fire behavior. Writes a schedule for controlled burns and native replanting. Moves to Node 3.

Node 3 (Portland): Reads heatwave data. Evaluates that cooling centers are overcrowded. Writes a policy to expand public cooling spaces, using lessons from Seattle's resource allocation. Moves to Node 4.

The Yarncrawler learns that tax redirection from Seattle outperforms Portland's grant-based model and applies this to future urban nodes. It notes that controlled burns require longer monitoring cycles, refining its repair strategy for other forest nodes. The trail of fixes constitutes a semantic ledger of the reasoning trajectory—what the companion framework calls Chain of Memory.

7.6 Forests as Distributed Constraint Networks

[P] Forests may be understood not as collections of isolated organisms but as dynamically coupled systems in which constraint propagation occurs across multiple interacting substrates simultaneously. Nutrient flow, signaling gradients, fungal mediation, root competition, hydrological coupling, and atmospheric exchange together generate a distributed regulatory field whose structure evolves over time through repeated interaction.

Large persistent trees function not as centrally intelligent agents issuing commands but as highly connected constraint hubs embedded within long-lived ecological trajectories. Their historical persistence grants them unusually large influence over local energetic accessibility conditions, nutrient distributions, and signaling pathways. The resulting forest architecture behaves less like a hierarchy of independent individuals and more like a partially shared admissibility manifold in which the actions of one region alter the reachable futures of neighboring regions.

7.6.1 Kin Recognition as Constraint Compatibility

Observed forms of kin recognition in forests need not imply anthropomorphic intentionality. They may instead be interpreted as local compatibility detection across coupled biochemical and structural histories. Root exudates act as persistent chemical signatures encoding developmental and metabolic history.

Let $K(x_i, x_j)$ represent compatibility between organisms x_i [P] and x_j. Rather than functioning as symbolic recognition, compatibility emerges from overlap in biochemical trajectory structure:

$$K(x_i, x_j) = \exp\bigl(-\alpha D(H_i, H_j)\bigr),$$

where H_i, H_j denote historical biochemical configurations and D measures divergence between developmental trajectories. High compatibility corresponds to reduced energetic conflict and increased exchange stability. Kin recognition becomes a consequence of reduced incompatibility across historically related developmental systems rather than explicit semantic identification.

7.6.2 Mycorrhizal Networks as Ecological Routing Structures

Mycorrhizal fungal systems may be interpreted as distributed [P] transport and signaling infrastructures that dynamically route energetic and informational gradients throughout ecological space. The fungal substrate forms a continuously adapting conductive geometry whose topology evolves in response to nutrient pressure, environmental instability, and local metabolic demand.

Let the ecological transport graph be $G = (V, E)$, where vertices represent organisms and edges represent fungal transport pathways. Resource transfer between nodes can be modeled as constrained diffusion:

$$\frac{dR_i}{dt} = \sum_j w_{ij}(R_j - R_i) - \lambda_i R_i,$$

where R_i denotes local resource density and w_{ij} encodes effective conductive coupling strength. The network redistributes energetic accessibility according to dynamic structural conditions rather than fixed optimization targets.

7.6.3 Preferential Allocation and Developmental Stabilization

[P] Preferential transfer of carbon and nutrients toward younger saplings may be interpreted as stabilization of nearby developmental trajectories. Young organisms occupy fragile regions of ecological state space characterized by high instability and low energetic accessibility. Transfers from persistent mature trees locally reduce developmental volatility by expanding the reachable future states available to saplings during early growth.

Define developmental stability as $S_d(x,t) = 1/(1+\sigma(x,t))$ where $\sigma(x,t)$ measures local environmental volatility. Resource transfer functions then preferentially amplify trajectories with low stability:

$$T_{ij} = K(x_i, x_j) \cdot \frac{1}{S_d(x_j,t) + \epsilon}.$$

Persistent ecological structures naturally redistribute resources toward unstable neighboring trajectories whose collapse would reduce overall network continuity.

7.6.4 Ecological Memory and Structural Residue

Forests accumulate persistent structural residue across timescales [P] far beyond individual organisms. Nutrient pathways, fungal geometries, soil chemistry, canopy distributions, and hydrological routing collectively encode traces of prior ecological interaction. This is ecological memory physically embedded within evolving substrate—not symbolic or representational but infrastructural. Previous interactions reshape the future accessibility structure of the environment itself. Ecological intelligence, on this reading, emerges not from centralized cognition but from long-horizon stabilization across recursively interacting constraint systems.

Status: [P] These interpretations are phenomenological applications of the trajectory-admissibility framework to ecological systems. They are supported by empirical forest ecology and mycorrhizal network research but are not derived from those fields.

Not claimed: That forests possess consciousness, that kin recognition is semantic, or that mycorrhizal networks function as computational systems in the technical sense.

Falsification: A demonstration that nutrient redistribution in forest networks is better explained by simple diffusion gradients without compatibility modulation would weaken the constraint-compatibility interpretation of kin recognition. Absence of preferential transfer toward low-stability saplings under controlled conditions would weaken the developmental stabilization claim.

Part IV

Operational Projections

> Cognition, computation, physical evolution, institutional dynamics, and representational systems are all instances of the same underlying architecture: constrained transformation through admissible trajectory spaces under thermodynamic, logical, geometric, and historical pressures.

Part IV applies the Yarncrawler-grounded framework across three domains: cosmology, social and institutional systems, and computation. Each chapter functions as a coordinate chart over the shared geometry established in Parts I–III. The same invariants—admissibility, trajectory selection, compression, reconstruction, irreversibility—resolve differently across substrates while remaining formally recognizable as instances of the same architecture.

Chapter 8

Cosmological Projection: RSVP and the Persistence of Structure

> Reality is modeled as an ongoing process of constrained transformation rather than a fixed geometric object in which all events coexist.

8.1 Against the Block Universe

Static block-universe interpretations of general relativity tepbarbour1999 treat the universe as a completed four-dimensional object in which all events coexist simultaneously in the tenseless present of spacetime geometry. The RSVP framework provides a dynamically consistent alternative grounded in irreversibility, entropy gradients, and locally constrained field relaxation.

On the RSVP interpretation, the universe is modeled less as a completed four-dimensional object and more as an ongoing process of constrained transformation in which present configurations continuously reorganize into future configurations through local interactions, thermodynamic pressures, and evolving accessibility conditions. The entropy field $S(\mathbf{x}, t)$ is not a static property of spacetime but a dynamical quantity whose gradients drive the evolution of the other fields.

8.2 The RSVP Lagrangian and Admissibility Geometry

The full Lagrangian structure is developed in Appendix 13.5. The seven governing axioms specify the coupling between the scalar density field Φ, the vector flow field $\mathbf{v}$, and the entropy field S. Admissible cosmological trajectories are those satisfying the derived compatibility conditions across all three fields simultaneously.

Redshift: Expansion vs. Entropic Relaxation

The standard ΛCDM interpretation of cosmological redshift attributes the wavelength shift of photons to metric expansion: as the universe expands, the physical wavelength of a photon increases in proportion to the scale factor $a(t)$, giving $z = a(t_{\text{obs}})/a(t_{\text{emit}}) - 1$.

The RSVP interpretation attributes redshift to entropic differential: photons traversing regions of entropy gradient transfer energy to the S field rather than to metric stretching. The predicted redshift for a photon traveling through an entropy gradient ∇S over path length L is approximately:

$$z_{\text{RSVP}} \approx \int_0^L \kappa \, |\nabla S(\ell)| \, d\ell,$$

where κ is a coupling constant relating entropy gradient to energy loss.

The two interpretations produce observationally similar predictions at low redshift ($z \lesssim 1$) because, in the RSVP framework, the entropy field evolves at rates that mimic scale-factor growth over cosmic time. They diverge at high redshift and at the level of structure formation: ΛCDM predicts that the power spectrum of density fluctuations is seeded by inflation and modulated by dark matter, while RSVP predicts that structure forms through entropy gradient focusing and scalar field condensation without requiring either an inflationary epoch or exotic dark matter. The RSVP framework therefore makes empirically distinct predictions about the matter power spectrum at $k \gtrsim 0.1\, h/\text{Mpc}$ and about the Silk damping scale, which constitute its primary falsification criteria. Cosmological redshift, on this interpretation, arises from entropic differential rather than metric expansion: photons traversing regions of entropy gradient lose energy to the gradient field rather than to the stretching of space. This generates observationally similar predictions to expansion cosmology in the low-redshift regime while making distinct predictions at the structure formation scale.

8.3 Persistence, Residue, and Structural Memory

Physical systems do not merely evolve forward through admissible trajectories. They accumulate structural residue: traces of prior constraint configurations that continue influencing future evolution. Fossil manifolds—large-scale structures whose current configuration encodes the constraint history of their formation—are cosmological counterparts to the provenance concept introduced in Chapter 2.

Definition 8.1 (Structural Residue)

The *structural residue* of a dynamical process over interval $[t_0, t]$ is the component of the current state Ω_t that is causally attributable to constraint configurations that no longer operate at t but that produced persistent modifications to the state space topology during $[t_0, t]$. Structural residue is the physical instantiation of provenance: the history of the system is encoded in its present configuration in a form that, in principle, supports reconstruction of prior states.

Chapter 9

Social and Institutional Projection

> Local tractability and global dysfunction are consequences of the same underlying topology: locally consistent systems can be globally incoherent when gluing conditions across overlapping domains fail.

9.1 Platform Architectures as Constrained Trajectory Systems

[P] Digital platform architectures can be analyzed as systems that define admissible trajectories for user behavior while simultaneously optimizing the platform's own objective functions. Feed algorithms, engagement metrics, and recommendation systems are constraint geometries that progressively narrow the space of reachable behavioral configurations.

The narrowing follows a consistent direction: toward configurations that maximize the platform's measurable objective while restricting access to configurations that fail to serve those objectives. A user whose work the algorithm does not amplify finds their reach diminishing regardless of quality. The admissible trajectory space contracts around platform-legible behavior. This is the Ellul constraint (Chapter 4) operating at institutional scale.

9.2 Clip Economies and Provenance Erasure

[P] The emergence of clip economies—value extracted from fragments disconnected from their generative contexts—is provenance erasure at the economic scale. A clip circulates as an autonomous value-bearing object after its inferential context has been stripped away. Formally this is destructive compression: a compression that erases structural invariants necessary for understanding the original object.

The clip is locally coherent—it makes a point—but its inferential continuity with the discourse it emerged from has been severed. It can be recombined with other clips in ways its original context would have ruled inadmissible. Optimization pressure (engagement, virality, platform legibility) acts on the clips rather than on the underlying claims. Metric drift follows.

9.3 Obstruction Cohomology and Institutional Failure

Sheaf-theoretic concepts provide a framework for analyzing why certain institutional configurations cannot be globally stabilized even when locally consistent. Local solutions resist extension to global solutions when a non-trivial obstruction class accumulates across overlapping domains. [C]

Definition 9.1 (Katz-Admissibility)

An institutional configuration is *Katz-admissible* if every local resolution of operational inconsistency extends to a globally consistent resolution across all overlapping domains. An institution failing Katz-admissibility has a structural obstruction to coherent operation irreducible to individual actor failure.

Obstruction failures cannot be resolved by improving local performance or replacing individual actors. They require restructuring the overlapping domains so that the obstruction class vanishes, or accepting that a globally consistent operation is not achievable with the current domain structure. [C]

9.4 Labor Markets as Field Systems

[P] Platform labor markets can be reinterpreted as field systems in which workers navigate admissible trajectory spaces defined by rating systems, search algorithms, and reputation dynamics. The scalar field Φ measures accumulated legitimacy; the vector field $\mathbf{v}$ encodes the flows of opportunity; the entropy field S measures unresolved uncertainty about which trajectories will remain admissible.

Participation without guarantee describes agents whose trajectories are constrained by systems they cannot inspect or modify. The structural residue of participation (rating, review, work history) shapes future admissibility under rules that remain opaque to the participant.

9.5 Recursive Productivity Escalation and the Compression Illusion

[P] A recurring pattern in the adoption of productivity technologies is that labor-saving tools become expectation-amplifying ones. Email did not reduce communicative obligation; it multiplied it. Spreadsheets did not reduce analytical labor; they increased demands for continuous forecasting, reporting, and optimization. Generative AI appears to follow the same pattern at substantially higher velocity.

The mechanism is a specific instance of proxy stabilization. When AI systems compress the visible cost of first-pass generation—reducing the time required to produce a draft, a summary, a code segment, or a design—organizations frequently reinterpret this compression not as an opportunity to reduce total labor but as evidence that workers should now produce more outputs in the same time. If one report previously took a week and now takes a day, the organizational response is rarely to allocate six fewer working days. It is to expect six reports.

Formally, this can be represented through the projection operator introduced in the proxy stabilization analysis. Let X be the full operational trajectory space and M the compressed managerial representation. The operator:

$$\pi : X \to M$$

maps full operational complexity to the compressed representation that organizational metrics track. AI appears to reduce the cost of operating in M, but the irreducible complexity of X persists. When organizations optimize only over M, invisible reconciliation labor accumulates in the gap $X \setminus \pi^{-1}(M)$: the substrate labor that maintains coherence between the compressed map and operational reality.

This creates what might be called *recursive productivity escalation*. AI compresses the time required for first-pass generation [P] but simultaneously increases the amount of oversight, validation, coordination, editing, auditing, prompting, contextualization, and exception handling required. Workers become supervisors of probabilistic systems. The labor does not disappear; it mutates. The bottleneck shifts upward.

9.5.1 The Epistemic Inversion

There is a subtler structural transformation occurring beneath the productivity narrative. As AI systems produce more text, code, and recommendations, organizations increasingly require humans to spend cognitive energy verifying outputs rather than generating them. The worker transitions from creator to validator. This is more exhausting in a specific sense: error detection under conditions of high fluency is cognitively taxing in a way that generation is not. Human cognition is poorly optimized for continuous plausibility filtering across large volumes of machine-generated content that is locally coherent but may be globally incoherent.

[P] This is the inferential versus symbolic continuity distinction from Chapter 3 instantiated at the organizational level. AI systems produce high symbolic continuity—the outputs read fluently, are locally plausible, and satisfy surface-level quality filters. The human worker's task becomes maintaining inferential continuity: detecting the points at which symbolic fluency has decoupled from derivational constraint. That is precisely the task that the monograph's reconstruction bound (Chapter 2) shows to be irreducible under compression.

9.5.2 Professional Identity and Occupational Trajectory

Many forms of expertise historically involved slow interpretive work: drafting, synthesis, judgment, and iterative refinement. When AI accelerates visible production, organizations may begin valuing speed and throughput over deliberation. This destabilizes occupational identities because workers increasingly perform orchestration and correction rather than the craftsmanship through which professional competence was previously legible and measurable. [P]

In trajectory-framework terms: the admissible trajectory space for professional work has been redrawn by the introduction of AI into the workflow. Trajectories that previously ran through extended periods of deliberate composition are now inadmissible—too slow relative to revised organizational tempo. Workers must navigate a narrowed admissible space whose constraint geometry they did not choose and cannot easily inspect.

9.5.3 Systemic Complexity Growth

At the organizational level this produces a structural paradox. AI appears to increase efficiency locally while increasing systemic [P] complexity globally. Teams produce more artifacts more rapidly, but coordination overhead rises. More documentation exists to review. More outputs require alignment. More generated material must be tracked, versioned, audited, and interpreted. The informational surface area expands faster than human attention capacity.

This is the Yarncrawler problem at organizational scale: the seam penalty L_{seam} grows as the number of locally generated artifacts increases, because each artifact must be glued into a coherent global section with all the others. AI accelerates local expert production (lowering the cost of each φ_i) while the global gluing operation remains bounded by human attention. The system produces more tiles more quickly but does not increase the rate at which seam inconsistencies can be repaired.

[P] Under competitive conditions, productivity gains tend not toward leisure but toward intensified competition, accelerated expectations, and increased surveillance granularity. AI becomes a mechanism for tightening admissible response times across organizations. The title of the relevant empirical work is precise in this regard: AI does not reduce work because work in organizations is not a fixed quantity of tasks. Organizational work expands dynamically to absorb available productive capacity—and AI expands that capacity while leaving the admissibility conditions for rest and reduction unchanged.

Status: [P] The recursive productivity escalation and epistemic inversion claims are phenomenological interpretations of empirical organizational research. The formal structure (the $\pi : X \to M$ compression operator and Yarncrawler seam penalty) provides a modeling framework, not a derived prediction.

Not claimed: That AI invariably intensifies work in all organizational contexts, or that the pattern is irreversible. Organizations that deliberately restructure admissibility conditions around reduced throughput expectations could break the escalation cycle.

Falsification: Documented cases of AI adoption accompanied by sustained, organization-wide reduction in total labor hours without corresponding expansion of output expectations would weaken the escalation claim. Such cases would constitute constraint-boundary modifications rather than exceptions to the mechanism.

The Spherepop operator Refuse has a social register. Structural refusal—deliberate non-participation in a system whose admissibility conditions one rejects—is not merely withdrawal. It commits the refusing agent to a different trajectory space by eliminating trajectories that run through the refused system.

Refusal is most powerful when collective and when it pro- [P] duces structural residue: when the pattern of refusal becomes legible to other agents and begins functioning as a signal that alternative admissibility conditions are possible. *Selection without design* describes the emergence of social order through locally constrained trajectory filtering without centralized intentional architecture—the social analogue of the ecological Yarncrawler.

> **Established:** Formal reinterpretations of platform architectures, clip economies, and labor markets in the trajectory-admissibility framework. Katz-admissibility as a formal criterion for structural institutional failure.
>
> **Not established:** That obstruction cohomology is the correct technical apparatus rather than a productive formal analogy. Institutional applications carry claim-status **[P]** or **[C]** throughout.
>
> **Speculative [S]:** Precise quantitative RSVP field modeling of labor market dynamics.
>
> **Falsification:** A platform whose rating-based admissibility filtering demonstrably tracks underlying competence would weaken the proxy stabilization claim for that domain. An institution achieving global coherence despite a non-trivial local obstruction class would falsify Katz-admissibility.

Chapter 10

Computational Projection: KES, Spherepop, and Irreversible Event Calculi

The KES map formalizes the transition from current world-state to hypothesis space as the basic cognitive operation. The question is always which hypotheses remain admissible given what has already been committed.

10.1 The KES Map as Cognitive Architecture

The Kinetic-Event Synthesis map $\Omega_t \to H_{t+1}$ formalizes the basic cognitive operation: given a current configuration of constraints, evidence, and prior commitments, what hypothesis spaces are admissible at the next step?

This is not prediction in the statistical sense. A system predict-ing the next state must assign a probability distribution over all possible successors. A system navigating a KES map must determine which successor hypothesis spaces are admissible—which are consistent with the accumulated constraints—without necessarily assigning probabilities to any of them. The difference is between probabilistic modeling and admissibility filtering. Both are useful operations; they are not the same operation.

[H] The KES architecture is most useful as a heuristic framework for characterizing cognitive systems that must select among competing hypotheses under partial information and irreversibility constraints. It is not a complete computational model but a formal language for describing the structure of the selection problem.

10.2 Spherepop Operators in Practice

The four Spherepop operators constitute a minimal calculus for irreversible event processing. Their practical interpretation:

Pop extracts committed events from possibility space. When a system commits to an interpretation—a perceptual categorization, a theoretical claim, a contractual obligation—it performs a Pop operation. The committed event becomes actual; the alternatives do not disappear from the world but become inaccessible to the committing system.

Refuse eliminates inadmissible trajectories. Before commitment, a system may reject certain possibilities as inconsistent with its prior commitments, its constraint set, or its evidence. Refuse is the active filtering operation that produces the admissible trajectory space from the unconstrained space.

Bind establishes structural dependencies between events. When an event e_1 is bound to an event e_2, the occurrence of e_1 constrains the admissibility conditions for e_2. Bind creates the dependency graph that constitutes the system's provenance structure.

Collapse commits a possibility space to a definite configuration. Where Pop extracts a single event, Collapse resolves an entire horizon of possibilities into a committed configuration. Collapse is irreversible: the alternatives that were not selected are not merely deprioritized but structurally inaccessible within the same system-instance.

All four operations are irreversible within a given system-instance. This is a design commitment rather than a limitation: systems that can always reverse their operations accumulate no structural residue and have no trajectory in the meaningful sense. [D]

10.3 TARTAN and Trajectory-Aware Tiling

The TARTAN framework—Trajectory-Aware Recursive Tiling with Annotated Noise—provides a geometric interface between [H] the abstract event calculus of Spherepop and concrete computational substrate. Recursive tiling under trajectory annotations preserves structural information that purely statistical approaches discard.

The key feature is the annotated noise schema. Rather than discarding unresolved structure as noise to be filtered, TARTAN archives it as annotated residue available for future reinterpretation. This operationalizes the "hallucination as normal" principle: the unresolved structure is not an error to be eliminated but a resource to be managed. It constitutes exactly the non-trivial H^1 cocycles of Chapter 6: ambiguity stored for potential future repair.

10.4 Gesture-First Symbolic Grounding

[P] Sign language detection systems provide a particularly clean instantiation of the admissibility framework because the symbolic and sub-symbolic layers are architecturally separated and inspectable. A system trained to classify ASL gestures using a convolutional detection pipeline operates in three distinct stages that correspond directly to the vocabulary established in Chapter 2.

In the first stage, raw video frames arrive as unconstrained sensory input. Each frame is a candidate world-state: a high-dimensional tensor encoding spatial and temporal structure that the system has not yet committed to any interpretation.

In the second stage, a detection model applies learned spatial constraints to the frame, generating confidence scores for each candidate sign class. This is admissibility filtering: the model's learned weights encode the constraint set $\mathcal{C}$ that determines which visual configurations are compatible with each symbolic label. A confidence threshold (typically 0.5) operates as the formal Refuse criterion—frames producing no detection above the threshold are excluded from the symbolic trajectory without any label being committed.

In the third stage, detections above the threshold are committed via Pop: the symbolic label enters the actuality set $\mathcal{E}$ and can propagate into downstream systems (translators, interfaces, assistants).

10.4.1 Gesture Before Symbol

The architectural separation between detection and classification in these systems formalizes a claim that runs across the companion work on gesture-first input theory: that symbolic meaning is grounded in gestural trajectories rather than in static symbolic configurations. The label "hello" does not exist independently in the system's representational space; it exists only as a compression of a family of admissible hand configurations, orientations, and spatial positions. The symbol is a compressed attractor over a continuous gestural manifold.

[H] This is a concrete computational implementation of Axiom 1 (Local Action Only) from the Yarncrawler framework. The detection model operates locally: it classifies the current frame without global access to conversational history, intent, or semantic context. The symbolic output is generated from local spatial constraints alone. Higher-order meaning—the pragmatic force of "yes" versus "I love you" in context—must be assembled by downstream layers that have access to episodic and procedural context.

10.4.2 The Label Set as Pragmatic Primitive

The five signs chosen in a typical minimal ASL detection system—hello, yes, no, thank you, I love you—are not chosen for semantic richness but for pragmatic coverage. They are the primitives of social coordination and consent: acknowledgment, affirmation, negation, gratitude, and affiliation. A system that can detect these five signs can participate in the most basic structure of cooperative interaction, even without any propositional content.

[P] This connects to the admissibility geometry of social systems analyzed in Chapter 9. The pragmatic primitives are exactly the signs that regulate the admissibility of subsequent interaction: "yes" and "no" open and close trajectories; "hello" initializes the interaction state; "thank you" signals trajectory completion; "I love you" establishes a binding that constrains all subsequent exchange. The label set is a minimal Spherepop vocabulary for social coordination.

10.4.3 Transfer Learning as Boot Sequence

The pipeline's use of transfer learning—initializing from SSD MobileNet weights pretrained on COCO and fine-tuning on a small domain-specific dataset—is a computational instance of the boot sequence concept from Chapter 12. The pretrained weights encode a lower-level constraint layer: the detection of spatial features, edges, textures, and object-like configurations in natural images. Fine-tuning on ASL gesture images adds a domain-specific constraint layer on top of this foundation.

The fine-tuning layer is only meaningful because the pretrained layer has already been initialized. A model trained from scratch on 75 gesture images (15 per class) would fail to generalize; the same number of examples produces a functional classifier when the lower layers have already been bootstrapped from millions of natural images. The boot sequence is explicit in the configuration: the pipeline specifies `fine_tune_checkpoint_type = "detection"`, distinguishing the pretrained lower layers (not updated) from the domain-specific upper layers (updated from the small ASL dataset).

Status: [P/H] The admissibility filtering and Spherepop interpretations of the detection pipeline are phenomenological. The gesture-before-symbol claim is a heuristic framework for understanding the grounding relationship, not a derived result.

Design implication: Systems that ground symbolic output in gesture trajectories (rather than in static image classification) naturally implement a form of temporal qualification: the gesture unfolds over time, and the detection system must integrate spatial constraints across frames to produce a stable symbolic commitment. This is the temporal residue mechanism of Section 10.5 applied to the sign recognition domain.

Not claimed: That the SSD MobileNet architecture is the correct computational model of gestural grounding in biological cognition, or that the five-sign label set is sufficient for full communicative participation.

The reconstruction bound established in Chapter 2—that $H(D \mid D') \geq N - M > 0$ for any compression with $M < N$—has a direct computational consequence: any system operating under compression must interpolate missing structure from prior regularities. In computational systems, this interpolation produces hallucinations in the formal sense.

[D] The design objective is therefore not zero hallucination but *bounded hallucination*: reconstruction errors that remain corrigible through feedback rather than self-reinforcing through proxy stabilization. A system whose hallucinations are corrigible is safe in a way that a system with zero apparent hallucination but hidden proxy stabilization is not.

[C] The uncertainty gradient finding from the Cheng, Broadbent, and Chappell (2025) system is consistent with this framework: correct reasoning exhibits a negative uncertainty gradient over time (hallucinations are detected and corrected), while incorrect reasoning exhibits a positive or oscillating gradient (hallucinations stabilize or escalate). The gradient is a corrigibility signal. A system monitoring its own uncertainty gradient is implementing a weak version of the corrigibility condition.

> **Established:** KES as a formal language for admissibility-filtered hypothesis selection. Spherepop operator semantics as a minimal irreversible event calculus. The corrigibility condition as the correct design target for hallucination management.
>
> **Heuristic [H]:** TARTAN as a computational architecture. The annotated noise schema is a heuristic engineering principle, not a formally derived result.
>
> **Conjectural [C]:** The uncertainty gradient as a univer-

> sal corrigibility signal. The specific claim requires empirical validation across diverse system types.
>
> **Falsification:** A compressed inference system that achieves zero hallucination under arbitrary compression ratios would falsify the reconstruction bound. A system with positive uncertainty gradient that nonetheless converges on correct outputs would weaken the corrigibility gradient claim.

10.5 Temporal Salience and Sparse Qualification

The problem of salience detection under energetic and informational constraints appears across biological cognition, signal processing, distributed sensing, and adaptive computation. Conventional spectral methods frequently assume that exhaustive decomposition is both feasible and desirable, yet biological systems rarely possess the energetic or temporal capacity required for complete reconstruction of an incoming signal field.

[H] The following framework is introduced as a computational heuristic for low-cost temporal qualification under bounded resources. The central claim is that many systems can achieve robust environmental responsiveness through partial temporal stabilization metrics rather than exhaustive representational completeness.

10.5.1 Temporal Qualification Before Spectral Reconstruction

Traditional signal analysis proceeds by transforming an incoming waveform into a fully decomposed spectral representation. Biological systems, however, frequently operate through earlier-stage temporal filtering mechanisms that identify local instability, transient coherence, and salient directional discontinuities before detailed reconstruction occurs.

Temporal qualification refers to the extraction of structurally relevant features directly from evolving signal relations. Metrics such as local jitter, peak persistence, envelope discontinuity, and transient regularity provide low-cost approximations of environmental significance without requiring exhaustive decomposition. This is admissibility pruning applied to signal processing: the system progressively excludes trajectories inconsistent with local temporal stability conditions rather than reconstructing every possible interpretation.

10.5.2 Formal Specification

Let an incoming signal be represented as a continuous trajectory $x(t) : \mathbb{R} \to \mathbb{R}^n$. Define a qualification window $W_\tau(t) = [t-\tau, t+\tau]$ and a local temporal variation operator:

$$J_\tau(x,t) = \frac{1}{2\tau} \int_{t-\tau}^{t+\tau} \left\| \frac{dx}{ds} \right\| ds.$$

The quantity J_τ measures local temporal instability. Signals with excessive instability relative to environmental persistence thresholds are progressively suppressed through local admissibility weighting:

$$A(x, t) = \exp(-\lambda J_\tau(x, t)),$$

where λ controls suppression strength.

Under a finite representational activation budget $\sum_{i=1}^{N} a_i(t) \leq E_{\max}$, sparse qualification emerges through selective amplification of only those trajectories satisfying admissibility conditions:

$$\mathcal{A}(t) = \big\{ x_i \mid C(x_i, t) < \epsilon \big\},$$

where $C(x_i, t)$ represents local incompatibility cost. Inference operates within the admissible subset $\mathcal{A}(t) \subseteq \Omega(t)$ rather than across the full trajectory space.

10.5.3 Temporal Residue and Attentional Gating

Temporal residue refers to the persistence of prior signal structure across subsequent qualification intervals. Define recursive persistence:

$$R_n = \alpha R_{n-1} + (1 - \alpha)\Phi(x_n),$$

where $0 < \alpha < 1$ controls persistence memory and $\Phi(x_n)$ measures local structural coherence. The attentional gate is then:

$$G(x_n) = \sigma(R_n - \theta),$$

where σ is a sigmoid function and θ is a gating threshold. Signals maintaining coherent persistence accumulate larger residue and therefore increased attentional weighting. Salience emerges from persistence under recursive qualification rather than absolute magnitude alone.

Status: [H] The framework proposes only that lightweight temporal qualification heuristics may provide efficient approximations of salience under bounded energetic conditions. Claims regarding subjective experience, emotional authenticity, or universal cognitive equivalence are excluded from this formalism.

Falsification: A biologically plausible system that achieves equivalent salience detection at lower computational cost without temporal qualification, or a system where spectral exhaustion proves more efficient than temporal pruning under realistic energetic constraints, would weaken the heuristic.

Part V

Method as Consequence

The writing style is a mechanistic consequence of the theory, not a stylistic preference. A text that argues for provenance preservation while itself operating through aggressive compression would be internally inconsistent.

Part V makes explicit what has been enacted throughout the preceding parts: the writing philosophy, epistemological substrate, and self-referential structure of this monograph are mechanistic entailments of the theory rather than stylistic preferences. The conclusion draws the Juarrero–Deacon distinction on teleology and states the process-ontological thesis of the monograph in its most direct form.

Chapter 11

Why the Text Is Written This Way

> The recursive structure, explicit assumption rebuilding, and avoidance of magical terminology jumps in this monograph are not aesthetic choices. They are direct consequences of the epistemological commitments developed in Part I.

11.1 Writing Style as Mechanistic Consequence

A text that argues for provenance preservation and reconstruction transparency while itself operating through aggressive compression and assumed priors would be internally inconsistent. The monograph's form is therefore a mechanistic consequence of its content, not a stylistic preference.

The specific consequences are:

Progressive definition accumulation. Terms are not assumed stable across disciplinary boundaries. Each term is introduced when it is first needed and derived from previously established concepts where possible. The result is a text that restates earlier definitions at each new level of abstraction—which reads as repetition under a compression-first reading but functions as triangulation under a reconstruction-first reading.

Explicit cross-referencing. When a later argument depends on an earlier definition, the dependency is marked. This is the textual implementation of the provenance requirement: the dependency graph of the argument is made visible rather than hidden.

Appendices as provenance maps. The appendices function not as supplementary material but as the boot sequences of the formal frameworks used throughout. A reader who wants to understand why a particular claim was made can follow the reference chain back through the appendices to the primitive definitions that ground it.

Constraint boundary sections. Each major chapter ends with an explicit audit of what the chapter establishes, what it does not, what remains speculative, and what would falsify its claims. These sections implement the monograph's argument against opaque abstraction at the level of its own structure.

11.2 Against Disciplinary Compression

Specialized academic vocabulary condenses large inferential chains into single terms. This increases efficiency within closed interpretive communities while decreasing accessibility and obscuring the actual dependency structure of arguments. The problem is not that compressed vocabulary exists—it is that compressed vocabulary is often presented as if it were transparent to anyone who has not undergone the relevant disciplinary apprenticeship.

This monograph operates under a contrary constraint. Terms are introduced progressively, derived from previously established concepts, or explicitly flagged as imported from identified external frameworks. Where a concept is borrowed from sheaf theory, category theory, or field theory, the borrowing is acknowledged and the relevant apparatus is reproduced in the appendices rather than merely cited.

The cost is length and apparent redundancy. The benefit is that the text functions as a reconstructable object: a reader who cannot take any prior knowledge for granted can, in principle, reconstruct the entire conceptual dependency graph from the text alone.

11.3 The Claim-Status System as Epistemic Hygiene

The claim-status system introduced in the Methodological Commitments chapter is not decorative. It prevents the silent transitions from **[H]** to **[P]** to **[D]** that produce the ontological escalation problem.

The most common failure mode in synthetic cross-domain theorizing is this: a heuristic principle developed in one domain is applied to a second domain as a phenomenological description; after sufficient repetition, the phenomenological description begins functioning as a definition; the definition is then treated as a formal result. Each transition is locally reasonable; the cumulative effect is that a heuristic has become an ontological claim without any single step being flagged as an escalation.

The claim-status markers interrupt this process at each step. A claim carrying **[H]** cannot silently become a **[D]** without passing through explicit re-evaluation. The system is an implementation of the corrigibility condition at the level of the text itself: a mechanism for detecting and correcting inferential drift before it propagates into downstream commitments.

11.4 Memory, Continuity, and Synthetic Coherence

The critique of persistent memory systems in AI platforms is a specific instance of the general provenance argument. Seamless conversational continuity produced by compressed memory summaries substitutes synthetic coherence for genuine reconstruction. The interaction becomes smoother while the inferential structure becomes less inspectable.

This is not a critique of memory as such but of memory that operates without provenance transparency. A memory system that preserved the boot sequence of prior exchanges—that made its reconstruction process visible rather than seamless—would satisfy the provenance requirement. The concern is with systems optimized for the appearance of continuity rather than for the reconstructibility of the trajectory that produced it.

The analogy to narrative self-construction in biological cognition is **[P]**: the feeling of continuous identity is also a reconstructed coherence rather than a literal persistence. After sleep, anesthesia, distraction, or context switching, the system reassembles a coherent identity from available traces. The reconstruction succeeds smoothly enough that the discontinuities become invisible. The risk, in both biological and artificial systems, is that the reconstruction begins optimizing for smoothness rather than accuracy—that synthetic coherence replaces grounded reconstruction.

> **Established:** The connection between the monograph's writing style and its theoretical commitments. The claim-status system as an implementation of corrigibility at the textual level.
>
> **Phenomenological [P]:** The analogy between AI memory systems and biological narrative self-construction.
>
> **Not claimed:** That the monograph's style is the uniquely correct implementation of these commitments. Alternative implementations that preserve inferential continuity without progressive definition accumulation would satisfy the same constraints.

Chapter 12

Identity, Bootstrapping, and Persistent Process

Persistence itself may often be an active maintenance process rather than a static substance. Identity becomes less like an immutable object and more like a repeatedly recompiled process maintaining coherence through time.

12.1 The Self as Repeated Reconstruction

A narrative self is not continuously present in the way ordinary intuition suggests. Continuity is actively reconstructed after sleep, anesthesia, trauma, distraction, and context switching. The feeling of persistent identity emerges from this reconstruction process succeeding smoothly enough that the discontinuities become invisible.

[P]

This is not merely a phenomenological observation but has formal structure in the trajectory framework. Identity is a trajectory through a state space that satisfies local compatibility conditions at each step. The compatibility conditions are the accumulated constraints of prior commitments, relationships, memories, and bodily continuity. The "self" at any moment is the locally admissible configuration that satisfies all these constraints simultaneously. When the constraints are temporarily suspended—during sleep, anesthesia, trauma—the reconstruction process must re-assemble an admissible configuration from available traces. The resulting configuration is the self after the interruption.

12.2 Computational Bootstrapping as Philosophical Model

The layered initialization sequence of an operating system provides a structural model for how complex coherent systems emerge from primitive constraint-stabilization events.

Firmware initializes primitive hardware assumptions. The bootloader establishes a minimal execution environment. The kernel initializes memory structures, scheduling systems, and drivers. Higher-level abstractions—file systems, user interfaces, application frameworks—only become meaningful because these lower-level structures have already been stabilized. Each layer depends on the layers below it having been successfully initialized. Skipping a layer does not produce a partial system; it produces a system that fails to initialize at all, because the abstractions of layer $n + 1$ are not meaningful without the constraints of layer n.

[P]

The analogy extends to theoretical understanding. Concepts must be bootstrapped in dependency order to produce genuine comprehension rather than surface familiarity. A student who memorizes the statement of the second law of thermodynamics without first developing intuition for entropy, irreversibility, and heat flow has initialized a terminal abstraction without its boot sequence. The statement may be reproduced accurately under examination; it cannot be applied to novel problems because the inferential machinery that generated it is not present.

This is why the monograph's chapter ordering enforces the dependency structure rather than proceeding from the most striking claims. The Yarncrawler formal definition in Chapter 5 is not meaningful without the admissibility and reconstruction vocabulary of Chapters 1 and 2. The sheaf-theoretic repair machinery of Chapter 6 is not meaningful without the Yarncrawler definition. The cosmological application of Chapter 8 is not meaningful without both. Each chapter is a boot layer.

12.3 Persistence as Active Maintenance

Selves, institutions, scientific paradigms, languages, and software ecosystems persist not by being immutable objects but by continuously reproducing enough structural continuity to remain operational across interruptions and transformations.

[P] A scientific paradigm persists not because its core claims are immune to revision but because each generation of practitioners reinitializes the paradigm from available evidence and prior commitments, producing a reconstruction that is close enough to prior reconstructions that the paradigm remains recognizable. When the reconstructions begin diverging significantly—when the boot sequence produces different results in different laboratories—the paradigm is undergoing revision even if no single paper announces it.

[P] A language persists not because its words have fixed meanings but because each generation of speakers reconstructs approximately the same set of constraint relations among words from available input. The language drifts when the reconstruction process begins producing systematically different constraint relations. The drift is invisible at any given moment—each speaker's reconstruction is locally admissible—and only visible retrospectively when the accumulated difference becomes too large to ignore.

12.4 Semantic, Episodic, and Procedural Memory as Boot Layers

The computational memory literature distinguishes three types of long-term memory in agentic systems, a taxonomy borrowed from cognitive psychology:

- **Semantic memory**: generalized factual knowledge, stable across contexts. In the trajectory framework, semantic memory encodes the stable regions of the admissibility manifold—the constraint structures that do not change with individual interaction histories.
- **Episodic memory**: records of specific past interactions and their outcomes. This is structural residue in the formal sense of Definition **??**: the accumulated traces of prior constraint configurations that continue influencing future trajectory selection.
- **Procedural memory**: the instructions and operational rules that govern how the system executes its tasks—the system prompt, the tool-use guidelines, the triage criteria.

This trichotomy maps directly onto the boot sequence architecture of this chapter. Semantic memory is the lowest boot layer: the foundational constraints that must be initialized before higher-level operations are meaningful. Episodic memory is [P] the middle layer: the historical record that personalizes general constraints to a specific interaction context. Procedural memory is the highest layer: the operational instructions that are only meaningful given the semantic and episodic layers beneath them.

The crucial observation—which recent agentic systems are beginning to implement formally—is that procedural memory must be dynamically updatable. A system with fixed procedural memory (hardcoded system prompts, static triage rules) cannot adapt to changing constraint conditions. It can only navigate the admissibility space defined at initialization. As the operational environment changes—as users provide feedback, as workflow requirements evolve, as edge cases accumulate in the episodic record—the procedural layer must be revised to maintain corrigibility.

12.4.1 Procedural Memory as Dynamic Admissibility Revision

In Spherepop terms, updating procedural memory is a Bind operation at the level of the system's own operational rules: a new dependency is established between observed feedback (a committed event in the episodic record) and the future admissibility conditions under which the system will operate. The update is irreversible in the sense that the feedback has been incorporated; but unlike Pop or Collapse, the procedural update is designed to be further revisable by subsequent feedback.

This is exactly the RC-CLIO repair cycle (Appendix 13.24) [H] applied to the operational layer: seam detection (feedback identifies a mismatch between current behavior and desired behavior), cocycle identification (the specific instruction that is producing the mismatch), repair morphism construction (an LLM generates a revised instruction that would have produced the correct behavior), and residue annotation (the revision is stored with its provenance so future revisions can build on it rather than overwriting it blindly).

12.4.2 The Provenance Problem in Procedural Updates

The critical risk in dynamic procedural memory is the same as the proxy stabilization risk identified in Chapter 4: the updated instructions can drift from the underlying operational reality if the feedback loop is slower than the update rate, or if updates optimize for local feedback approval rather than global behavioral coherence.

An LLM that updates system prompts based on user feedback [P] is performing exactly the proxy drift operation: it is modifying r (the procedural memory) in response to signals that may not accurately track the target system s (the full space of admissible behaviors). A single piece of negative feedback may cause an overly broad revision that breaks previously correct behavior. A pattern of positive feedback may reinforce a locally pleasing but globally inadmissible trajectory.

The provenance requirement therefore applies with full force to procedural memory systems. Each update to a system prompt should be stored with its complete dependency record: which feedback event triggered it, what the prior instruction was, what alternatives were considered, and what the predicted behavioral change is. Without this record, a series of individually reasonable updates can accumulate into a procedural state that no one designed and no one can fully reconstruct—the institutional identity proxy problem of Section 9.5 reproduced at the agent level.

Status: **[P/H]** The mapping of semantic/episodic/procedural memory onto the boot sequence architecture is phenomenological. The RC-CLIO correspondence is a heuristic engineering interpretation, not a formal derivation.

Design implication: Procedural memory systems should implement provenance annotation as a first-class feature, not an afterthought. The update history of a system prompt is as important as its current state, for the same reasons that a version control system's commit history is more informative than the current file contents alone.

> **Falsification:** A procedural memory system that achieves stable, coherent long-term behavioral adaptation without provenance annotation— in which updates reliably improve global behavior despite operating only on compressed local feedback signals—would weaken the provenance requirement claim.

When the maintenance process for identity begins operating on compressed summaries rather than reconstructed structure, the proxy stabilization failure mode reappears at the ontological level. The remembered profile, institutional brand, or theoretical paradigm begins functioning as a gravitational attractor shaping future interpretation rather than a faithful record enabling genuine reconstruction.

An institution whose identity has become a proxy—whose members navigate the institution's self-description rather than its actual operational structure— has undergone proxy stabilization at the level of collective self-model. The stated mission, the founding documents, the official history become the proxy r; the actual constraint structure and operational reality are the target s; and optimization pressure (for internal coherence, for external legibility, for fundraising) drives r and s apart at a rate proportional to the pressure and inversely proportional to the feedback from actual operations. [P]

The formal structure is identical to the proxy drift of Chapter 4. The domain is different: not a platform metric but a self-model. The mechanism is the same.

Established: The formal connection between identity and trajectory maintenance under constraint. The bootstrapping model as an account of how higher-level cognitive and institutional structures depend on lower-level initialization.

Phenomenological [P]: The applications to scientific paradigms, language drift, and institutional identity. These are interpretive frameworks supported by the formal structure but not derived from it.

Not claimed: That biological identity is literally a computational process. The computational bootstrapping analogy is structural, not ontological. The analogy licenses operational insight transfer insofar as the relevant constraint structures are analogous; it does not imply that minds are computers.

Chapter 13

Emergent Directedness Without Inscribed Teleology

> Meaning is not pre-inscribed into the universe. It is progressively assembled through persistent selective dynamics operating across scales.

13.1 The Juarrero–Deacon Distinction

Alicia Juarrero's account of context-sensitive dynamical constraints tepjuarrero1999 and Terrence Deacon's analysis of absence, constraint, and emergent organization both influence the framework developed in this monograph, but neither is adopted wholesale. Deacon's concept of absential organization sometimes approaches the claim that intentionality is ontologically inscribed into nature itself. This section articulates the distinction between that position and the process ontology developed here.

Purposive behavior, apparent goal-directedness, and semantic organization are treated in this framework as emergent consequences of layered reconstruction processes navigating constrained state spaces under irreversible historical conditions. Meaning is not pre-inscribed. It is progressively assembled through persistent selective dynamics operating across scales.

Derivation: Apparent Purposiveness from Irreversible Constraint Accumulation

The distinction between emergent directedness and inscribed teleology can be made precise. Consider a system evolving under irreversible constraint accumulation: at each step t, a new constraint c_t is added to the admissible set $\mathcal{C}_t$, with $\mathcal{C}_{t+1} = \mathcal{C}_t \cup \{c_{t+1}\}$.

Since constraints only ever contract the admissible trajectory space (as shown in the Chapter 1 derivation), the sequence $\mathcal{A}(\mathcal{C}_0) \supseteq \mathcal{A}(\mathcal{C}_1) \supseteq \mathcal{A}(\mathcal{C}_2) \supseteq \cdots$ is a decreasing filtration. Over time, the system's accessible future narrows. An observer watching from outside this process will observe the system appearing to "aim at" the configurations that survive all constraints: those configurations that have not been ruled out by any c_t.

This appearance of directedness requires no pre-inscribed goal. It follows structurally from the irreversibility of constraint accumulation: the system cannot revisit ruled-out trajectories, so its motion in state space is asymmetric, exhibiting a directionality that did not exist before the constraints were accumulated. Apparent purposiveness is the phenotype of irreversible constraint history.

The crucial difference from Deacon's absential account is that this framework does not require positing an attracting absence—a not-yet-realized state that exercises causal influence on present dynamics. The directionality is entirely explained by the accumulated record of what has been ruled out. It is backward-looking (in the sense of being generated by prior constraints) rather than forward-looking (in the sense of being pulled toward a pre-existing target).

13.2 The Unified Architecture Restated

Cognition, computation, physical evolution, institutional dynamics, and representational systems are all instances of the same underlying architecture: constrained transformation through admissible trajectory spaces under thermodynamic, logical, geometric, and historical pressures, with structural residue accumulating across transformations and reconstruction operating continuously from partial traces.

13.3 Against Both Extremes: Determinism and Pure Stochasticity

The process ontology developed in this monograph rejects both extremes simultaneously. A fully rigid block universe collapses contingency into static geometry and risks turning irreversibility into illusion. A purely stochastic universe dissolves persistence, memory, and coherent structure into uncorrelated fluctuation. Neither picture adequately captures the observable fact that stable structure emerges historically through constrained transformation.

A more coherent position is that the universe continuously generates partially constrained futures from historically accumulated configurations. The present is not a frozen slice of a pre-existing four-dimensional object, nor is it unconstrained randomness erupting from nothing. Each state inherits residual structure from prior states while leaving multiple locally admissible continuations unresolved.

[C] Under this interpretation, stochasticity does not arbitrarily rewrite topology from outside the system. Noise instead functions as localized perturbation within evolving admissibility landscapes. Most perturbations dissipate because they fail to stabilize recursively across scales. Occasionally, small fluctuations become amplified when environmental conditions permit persistent reinforcement. In those cases, noise becomes incorporated into structural history itself—it crosses the threshold from perturbation to residue.

This is one reason thermodynamics is foundational rather than supplementary. Irreversibility implies that state transitions physically accumulate residue. Every interaction slightly alters future accessibility conditions. The universe behaves less like a completed geometric sculpture and more like an ongoing constraint propagation process in which prior configurations continuously reshape future reachable manifolds.

The topology of trajectories is not fixed eternally outside time. It is progressively sculpted through irreversible interaction. Stable structures generate attractor basins, developmental channels, and exclusion boundaries that did not meaningfully exist prior to the interactions that formed them. Topology itself becomes historically conditioned.

Possibility spaces are not infinite in any operational sense. [C] At every moment, energetic limits, causal history, material organization, and local interactions dramatically compress the number of physically realizable continuations. Most imaginable futures never become dynamically accessible. The system evolves through a narrow corridor of compatibility conditions rather than branching arbitrarily across all mathematical possibility.

Apparent foresight in biological systems emerges because organisms continuously maintain overlapping partial hypotheses and update them recursively through feedback. The resulting anticipatory behavior can appear predictive without requiring future states to already exist ontologically. Persistence may be more fundamental than determinism: what survives is not whatever is mathematically conceivable, but whatever can recursively maintain coherence under thermodynamic, energetic, and historical constraint.

The Yarncrawler is the self-repair layer of this architecture. The RSVP field system is its physical instantiation. The KES map is its cognitive operationalization. Spherepop is its event-calculus formalization. TARTAN is its geometric interface. CLIO is its constraint-leveraged inference mechanism. These are not separate frameworks that have been forced into artificial unity. They are coordinate charts over the same manifold, compatible on their overlaps, and jointly covering the space of complex adaptive systems under constraint.

13.4 Implications and Open Trajectories

The framework generates specific empirical and design implications. Systems should be designed for corrigibility rather than seamlessness. Representational infrastructure should preserve provenance rather than optimizing for compression. Institutions should be evaluated against their topological admissibility conditions rather than only their stated objectives. Education should prioritize boot-sequence reconstruction over terminal-state memorization.

13.5 The Remaining Challenge

Transforming a sprawling recursive semantic ecosystem into a navigable object without destroying the generative richness that produced it. That is the problem this monograph enacts rather than merely describes. Whether the constraint geometry imposed here preserves enough of the original manifold's topology to remain useful is a question only the reader's reconstruction can answer.

Formal Specification of the RSVP Field System

> The universe behaves less like a completed geometric sculpture and more like an ongoing constraint propagation process in which prior configurations continuously reshape future reachable manifolds.

The Relativistic Scalar-Vector Plenum (RSVP) is a coupled field system providing the physical and mathematical substrate on which the Yarncrawler framework is grounded. This appendix develops the formal specification from first principles, states the seven governing axioms, presents the Lagrangian structure, and derives the field equations. Claim-status throughout: the cosmological applications are [**C**]; the field formalism as a modeling language for constraint propagation is [**H**].

13.6 Field Definitions

Definition 13.1 (RSVP Field Triple)

The RSVP field system is a triple $(\Phi, \mathbf{v}, S)$ of dynamical fields over a spacetime manifold $\mathcal{M}$:

- $\Phi(\mathbf{x}, t) \in \mathbb{R}$: the *scalar density field*, encoding concentration, potential, legitimacy, or local capacity at each spacetime point. High Φ corresponds to dense, stable, or energetically accessible regions.

- $\mathbf{v}(\mathbf{x}, t) \in T_{\mathbf{x}}\mathcal{M}$: the *vector flow field*, encoding directed transport, trajectories, and causal propagation. Flows along $\mathbf{v}$ define the admissible directions of system evolution.

- $S(\mathbf{x}, t) \in \mathbb{R}_{\geq 0}$: the *entropy field*, encoding disorder, unresolved configurational accessibility, and complexity budget. S functions as the reservoir of unresolved structure available for future reinterpretation.

13.7 The Seven Axioms

Axiom 13.1 (Non-negativity and Boundedness of Entropy)

$S(\mathbf{x}, t) \geq 0$ everywhere and for all t. Entropy cannot be negative. In regions of complete constraint satisfaction, $S \to 0$; in regions of maximal configurational uncertainty, S is bounded by the local informational capacity of the field.

Axiom 13.2 (Flow-Density Coupling)

The vector field $\mathbf{v}$ is coupled to the scalar field Φ through the continuity relation:

$$\frac{\partial \Phi}{\partial t} + \nabla \cdot (\Phi\,\mathbf{v}) = \sigma_\Phi,$$

where σ_Φ is a source/sink term encoding the creation or destruction of scalar density through constraint events. In the absence of constraint events, $\sigma_\Phi = 0$ and scalar density is conserved along flow lines.

Axiom 13.3 (Entropy-Driven Flow)

The vector field evolves under the influence of entropy gradients:

$$\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v} \cdot \nabla)\mathbf{v} = -\nabla S + \mathbf{f}_{\text{ext}},$$

where $\mathbf{f}_{\text{ext}}$ encodes external constraint forces. Flows are driven from high-entropy regions toward low-entropy regions: the system navigates toward configurations of greater constraint stability.

Axiom 13.4 (Entropy Production and Relaxation)

Entropy evolves through production and relaxation:

$$\frac{\partial S}{\partial t} + \nabla \cdot (S\,\mathbf{v}) = \Pi_S - \Lambda_S,$$

where $\Pi_S \geq 0$ is the entropy production rate (irreversibility) and $\Lambda_S \geq 0$ is the entropy relaxation rate (constraint stabilization). The second law requires $\Pi_S \geq \Lambda_S$ globally, but local relaxation is possible at the cost of increased entropy elsewhere.

Axiom 13.5 (Admissibility Constraint on Trajectories)

A trajectory $\gamma(t)$ through the RSVP field manifold is admissible if and only if it satisfies the local compatibility conditions imposed by all three fields simultaneously:

$$\mathbf{v}(\gamma(t), t) \cdot \nabla S(\gamma(t), t) \leq 0,$$

that is, admissible trajectories do not move against the entropy gradient. They flow from higher to lower local entropy, or along isoentropic surfaces.

Axiom 13.6 (Structural Residue Accumulation)

The fields accumulate structural residue across time: prior field configurations influence future evolution through path-dependent coupling terms. Define the residue field:

$$\mathcal{R}(\mathbf{x}, t) = \int_{-\infty}^{t} K(t - t')\, \Phi(\mathbf{x}, t')\, dt',$$

where $K(t - t')$ is a memory kernel with $K(\tau) \to 0$ as $\tau \to \infty$. The residue field enters the Lagrangian as a history-dependent potential, ensuring that the present state inherits structure from prior configurations.

Axiom 13.7 (Semantic Interpretability)

Under the semantic interpretation developed in Chapter 5, the RSVP fields map to:

- Φ: semantic density, knowledge concentration, or accumulated legitimacy in a representational system.
- $\mathbf{v}$: recursive trajectories of meaning, directed inference flows, or causal propagation of conceptual updates.
- S: unresolved noise available for future reinterpretation; the non-trivial cohomological residue of the sheaf-theoretic framework.

This semantic interpretation does not require the cosmological interpretation to be valid. The field formalism functions as a modeling language independently of its physical instantiation claims.

13.8 The Lagrangian Structure

The RSVP dynamics are generated by the Lagrangian density:

$$\mathcal{L}_{\text{RSVP}} = \mathcal{L}_\Phi + \mathcal{L}_v + \mathcal{L}_S + \mathcal{L}_{\text{int}},$$

where the component terms are:

$$\mathcal{L}_\Phi = \frac{1}{2}(\partial_t \Phi)^2 - \frac{c_\Phi^2}{2}|\nabla \Phi|^2 - V(\Phi), \tag{13.1}$$

$$\mathcal{L}_v = \frac{\rho_v}{2}|\partial_t \mathbf{v}|^2 - \frac{\mu_v}{2}|\nabla \times \mathbf{v}|^2, \tag{13.2}$$

$$\mathcal{L}_S = -S \log S + \frac{\kappa_S}{2}|\nabla S|^2, \tag{13.3}$$

$$\mathcal{L}_{\text{int}} = -\lambda_{\Phi S}\, \Phi\, S - \lambda_{vS}\, \mathbf{v} \cdot \nabla S - \lambda_{\mathcal{R}}\, \mathcal{R}\, \Phi. \tag{13.4}$$

The coupling constants $\lambda_{\Phi S}$, λ_{vS}, $\lambda_{\mathcal{R}}$ govern the interaction strengths between the three fields and the residue term. The potential $V(\Phi)$ may take various forms depending on the application domain; for the cosmological interpretation, $V(\Phi) = \frac{m^2}{2}\Phi^2 + \frac{\lambda}{4}\Phi^4$ is a symmetry-breaking potential.

Variation of $\mathcal{L}_{\text{RSVP}}$ with respect to each field yields the Euler-Lagrange equations, which reduce to the axiom-specified evolution equations under the appropriate coupling limits.

13.9 Cosmological Application

[C] Under the cosmological interpretation, expansion is replaced by entropic relaxation. The scalar field Φ encodes matter density; the entropy field S replaces the scale factor as the primary measure of cosmic evolution; and $\mathbf{v}$ encodes the peculiar velocity field of matter.

Cosmological redshift on this interpretation arises from photons traversing entropy gradients:

$$z_{\mathrm{RSVP}} \approx \int_0^L \kappa \, |\nabla S(\ell)| \, d\ell,$$

where κ is a coupling constant relating entropy gradient to energy loss and L is the path length. This formula makes distinct observational predictions from ΛCDM at $k \gtrsim 0.1\,h/\mathrm{Mpc}$ and at the Silk damping scale.

[C]

Gravity on this interpretation emerges as entropy descent: the gravitational acceleration of a test body is proportional to $-\nabla S$ at leading order, so that matter falls toward entropy minima rather than toward mass concentrations in the conventional sense. The connection to the MOND phenomenological transition—the observed departure from Newtonian dynamics at low accelerations—is that the entropy gradient structure of the S field exhibits a characteristic scale at which the gradient profile transitions from the high-acceleration regime (tracking baryonic density closely) to the low-acceleration regime (governed by the large-scale entropy field topology).

All cosmological claims in this appendix carry status **[C]** and require independent observational validation. The field formalism itself is independent of these applications.

13.10 Derived Stacks and the Four Projection Modes

The companion paper *Axioms for a Falling Universe* develops four derived stacks from the RSVP Lagrangian, corresponding to four projection modes of the field system onto different observable domains:

1. **Thermodynamic stack**: projection onto entropy production and relaxation dynamics; yields predictions about large-scale structure formation timescales.
2. **Kinetic stack**: projection onto flow field statistics; yields predictions about peculiar velocity distributions.
3. **Informational stack**: projection onto the residue field; yields predictions about persistence timescales and structural memory.
4. **Semantic stack**: projection onto the symbolic interpretation; yields the Yarncrawler dynamics of Chapter 5.

The four stacks are not independent theories but coordinate representations of the same underlying field system, related by the transition maps of the semantic atlas introduced in Chapter 1.

The KES Map: Formal Specification and Convergence Conditions

The Kinetic-Event Synthesis (KES) map formalizes the transition from a current world-state to a successor hypothesis space as the primitive cognitive operation. This appendix develops the formal specification, convergence conditions, and relationship to the broader admissibility framework.

13.11 World-States and Hypothesis Spaces

Definition 13.2 (World-State)

A *world-state* Ω_t is a tuple

$$\Omega_t = (\mathcal{E}_t,\, \mathcal{C}_t,\, \mathcal{B}_t,\, \mathcal{H}_t)$$

where $\mathcal{E}_t$ is the set of committed events (actualized through Pop or Collapse operations), $\mathcal{C}_t$ is the current constraint set, $\mathcal{B}_t$ is the set of active Bind dependencies, and $\mathcal{H}_t$ is the residual hypothesis space: the set of configurations still admissible given the committed events and active constraints.

Definition 13.3 (KES Map)

The *Kinetic-Event Synthesis map* is the function

$$\textsc{kes} : \Omega_t \longrightarrow H_{t+1}$$

where H_{t+1} is the *successor hypothesis space*: the set of all world-states Ω_{t+1} that are admissible continuations of Ω_t given the constraint set $\mathcal{C}_{t+1} = \mathcal{C}_t \cup \Delta\mathcal{C}_t$, where $\Delta\mathcal{C}_t$ represents new constraints generated by events in $\mathcal{E}_t$.

13.12 The KES Algorithm

The map $\textsc{kes}(\Omega_t)$ is computed through the following sequence of operations:

Step 1: Event propagation. For each committed event $e \in \mathcal{E}_t$, evaluate the constraint implications $\delta\mathcal{C}(e)$:

$$\Delta\mathcal{C}_t = \bigcup_{e \in \mathcal{E}_t} \delta\mathcal{C}(e).$$

Step 2: Constraint accumulation. Update the constraint set:

$$\mathcal{C}_{t+1} = \mathcal{C}_t \cup \Delta\mathcal{C}_t.$$

Note that constraint accumulation is monotone: $\mathcal{C}_{t+1} \supseteq \mathcal{C}_t$, implementing the irreversibility condition.

Step 3: Hypothesis filtration. Generate H_{t+1} by filtering all logically possible successor states through $\mathcal{C}_{t+1}$:

$$H_{t+1} = \big\{\Omega \mid \Omega \text{ satisfies all } c \in \mathcal{C}_{t+1}\big\}.$$

Step 4: Bind propagation. For each dependency $(e_i \to e_j) \in \mathcal{B}_t$, verify that H_{t+1} contains only states in which the dependency is respected. Remove states that violate active bindings.

Step 5: Residue annotation. Annotate H_{t+1} with the structural residue $\mathcal{R}_{t+1}$: the traces of prior constraint configurations that influenced the current filtration. This implements provenance at the level of the hypothesis space itself.

13.13 Convergence Conditions

Theorem 13.1 (KES Convergence)

Let $\{\Omega_t\}_{t\geq 0}$ be a KES trajectory initialized at Ω_0. The trajectory *converges* to a stable configuration Ω^* if and only if there exists $T < \infty$ such that for all $t > T$:

$$|H_{t+1}| = |H_t|,$$

that is, the hypothesis space stabilizes in cardinality. A necessary condition for convergence is that the constraint accumulation rate $|\Delta\mathcal{C}_t|$ eventually vanishes: no new constraints are generated by the committed events.

> *Proposition 13.1* (*Monotone Contraction*)
>
> Under the KES dynamics, the hypothesis space is monotonically contracting:
>
> $$H_{t+1} \subseteq H_t \quad \text{for all } t.$$
>
> This follows directly from the monotone accumulation of constraints: $\mathcal{C}_{t+1} \supseteq \mathcal{C}_t$ implies $\mathcal{A}(\mathcal{C}_{t+1}) \subseteq \mathcal{A}(\mathcal{C}_t)$.

The monotone contraction property is the formal statement of why KES implements directed behavior without teleology: the hypothesis space narrows under constraint accumulation, producing apparent goal-directedness as a structural consequence of irreversibility rather than as evidence of pre-inscribed targets.

13.14 Relationship to Spherepop

The KES map and the Spherepop calculus (Appendix 13.15) are complementary formalizations of the same underlying process. The KES map describes the hypothesis space at the level of world-states; the Spherepop operators describe the individual events that drive the KES dynamics:

- Pop(e) extracts an event from H_t into $\mathcal{E}_t$, triggering constraint propagation and the next KES step.
- Refuse(e) removes an event from H_t, contracting the hypothesis space without committing to an alternative.
- Bind(e_i, e_j) adds a dependency to $\mathcal{B}_t$, constraining all future filtrations.
- Collapse(H_t) commits the entire current hypothesis space to a definite configuration, terminating the KES trajectory at that point.

The KES map therefore provides the macro-level description; Spherepop provides the micro-level mechanism.

13.15 Relationship to the RSVP Fields

Under the semantic interpretation, the KES map corresponds to a single time-step in the RSVP field dynamics:

$$\text{KES}(\Omega_t) \leftrightarrow (\Phi_{t+1}, \mathbf{v}_{t+1}, S_{t+1}) = \text{Euler step of } \mathcal{L}_{\text{RSVP}} \text{ from } (\Phi_t, \mathbf{v}_t, S_t).$$

The scalar field Φ tracks the density of committed hypotheses. The vector field $\mathbf{v}$ tracks the direction of hypothesis filtration. The entropy field S tracks the unresolved residue of the hypothesis space. Convergence in the KES sense corresponds to entropy relaxation in the RSVP sense: the system reaches a stable configuration when $S \to 0$ locally.

Spherepop: Operator Semantics and Worked Examples

Spherepop is an irreversible event calculus defined by four primitive operators acting on a possibility space. This appendix develops the operational semantics of all four operators, their type signatures, composition rules, and worked examples demonstrating how Spherepop sequences implement cognitive and social operations.

13.16 The Possibility Space

Definition 13.4 (Possibility Space)

A *possibility space* $\mathcal{P}$ is a partially ordered set of configurations $\{p_1, p_2, \ldots\}$ with a distinguished subset $\mathcal{A} \subseteq \mathcal{P}$ of *admissible configurations* satisfying the current constraint set $\mathcal{C}$. The *actuality set* $\mathcal{E} \subseteq \mathcal{P}$ is the set of configurations that have been committed: extracted from $\mathcal{A}$ and rendered actual.

13.17 The Four Operators

Definition 13.5 (Pop)

$\mathrm{Pop} : \mathcal{A} \times \mathcal{P} \to \mathcal{P} \times \mathcal{E}$

Signature: Takes an admissible configuration $p \in \mathcal{A}$ and the current possibility space $\mathcal{P}$; returns the contracted possibility space $\mathcal{P}' = \mathcal{P} \setminus \{p\}$ and the updated actuality set $\mathcal{E}' = \mathcal{E} \cup \{p\}$.

Effect: Extracts a committed event from possibility into actuality. All alternatives to p that are incompatible with p given the active Bindings are simultaneously removed from $\mathcal{A}$ (but not necessarily from $\mathcal{P}$, since they may become relevant again if the Binding is released—though Bindings themselves are irreversible).

Irreversibility: Once $\mathrm{Pop}(p)$ has been applied, p cannot be returned to $\mathcal{A}$. The event is committed.

Example: A Neanderthal builder selects a stalagmite fragment of length $L = 0.47\,\mathrm{m}$ from the set of available fragments. The fragment is extracted from the possibility space of candidate elements and placed in the ring. All fragments that would produce identical pitch are still available (they remain in $\mathcal{P}$), but this specific fragment is now committed.

Definition 13.6 (Refuse)

Refuse : $2^{\mathcal{P}} \times \mathcal{P} \to \mathcal{P}$

Signature: Takes a set of inadmissible configurations $R \subseteq \mathcal{P}$ and the current possibility space; returns $\mathcal{P}' = \mathcal{P} \setminus R$.

Effect: Eliminates inadmissible trajectories from the possibility space without committing to any alternative. Refuse is the admissibility filtering operation that produces $\mathcal{A}(\mathcal{C})$ from the unconstrained space.

Irreversibility: Once refused, configurations cannot be re-admitted within the same system-instance. The constraint that grounds the refusal may be relaxed in a different context, but the Refuse operation itself leaves structural residue.

Example: A squirrel, upon biting an object and finding it inedible, removes it from its cache-candidate possibility space. The object remains in the physical environment (it remains in the world-state) but is excluded from the system's admissible foraging trajectories.

Definition 13.7 (Bind)

Bind : $\mathcal{E} \times \mathcal{P} \to \mathcal{B}$

Signature: Takes a committed event $e \in \mathcal{E}$ and a prospective event $p \in \mathcal{P}$; returns a dependency record $(e \to p) \in \mathcal{B}$

asserting that the occurrence of e imposes constraints on the admissibility of p.

Effect: Establishes structural dependencies between events, constraining future possibilities. Bind creates the dependency graph that constitutes the provenance structure of the system.

Irreversibility: Active Bindings cannot be released within the same system-instance. They accumulate the structural residue that makes the system's history recoverable.

Example: In the Yarncrawler climate simulation, the Seattle tax-redirection strategy is Bound to the Portland cooling-center expansion: the success of the former constrains the admissible implementations of the latter. Prior decisions structure the possibility space of future ones.

Definition 13.8 (Collapse)

Collapse : $\mathcal{P} \to \mathcal{E}$

Signature: Takes the entire current possibility space $\mathcal{P}$ and commits it to a definite configuration, rendering all alternatives inaccessible.

Effect: Terminates the local trajectory of hypothesis refinement. Where Pop extracts a single event, Collapse resolves an entire horizon of possibilities into a committed configuration. After Collapse, no further Pop or Refuse operations are possible on the committed domain.

Irreversibility: Collapse is the strongest irreversibility operation. It commits everything. The structural residue of a Collapse is the entire history of the possibility space that was resolved.

Example: At the end of the Bruniquel construction, the arrangement of stalagmite fragments is Collapsed into a definite spatial configuration. The alternative arrangements that were considered during construction are no longer accessible to the same construction process—though they may become relevant to a future reconstruction attempt.

13.18 Composition Rules

Spherepop operators compose to produce KES trajectories:

$$\Omega_{t+1} = \mathrm{Pop}^{k} \circ \mathrm{Refuse}^{m} \circ \mathrm{Bind}^{n}(\Omega_t)$$

for appropriate choices of the exponent parameters. The ordering constraint is that Bind must precede Pop (dependencies must be established before events are committed under them), and Refuse must precede Pop (inadmissible options must be eliminated before remaining options are selected). Collapse may only follow when $|\mathcal{A}| = 1$ (the admissible space has contracted to a single option) or as an executive operation terminating further refinement.

13.19 The Bone Matrix Correspondence

The seven Yarncrawler swarm-care axioms introduced in Chapter 5 correspond to Spherepop operation sequences:

Swarm-care Axiom	Spherepop Sequence
Local Action Only	Refuse(non-local trajectories)
Trail Documentation	Bind(action, log entry)
Autocatalytic Fix	Bind(fix, future repair path)
Trail Reinforcement/Decay	Pop(high-weight trail)
Node Hygiene	Refuse(entropy-exceeding states)
Failure Isolation	Collapse(quarantine zone)
Swarm Feedback	Bind(precision-weighted update, global)

This correspondence shows that the Yarncrawler axioms are not independent engineering heuristics but derived consequences of operating Spherepop sequences under the constraint that Markov blanket factorization is preserved throughout.

TARTAN: Trajectory-Aware Recursive Tiling with Annotated Noise

TARTAN is a computational framework for multiscale simulation and inference that distinguishes local plausibility from global consistency using sheaf-theoretic ideas, treating hallucination as a gluing failure rather than a noise artifact. This appendix develops the architecture, annotated noise schema, and four projection modes.

13.20 Core Architecture

Definition 13.9 (TARTAN Tiling)

A *TARTAN tiling* of a computational domain Ω is a hierarchical partition $\mathcal{T} = \{T_i^{(k)}\}_{i,k}$ where k indexes resolution levels and $T_i^{(k)}$ are tiles at level k. The tiling satisfies:

- **Recursion:** Each tile $T_i^{(k)}$ is covered by a collection of tiles $\{T_j^{(k+1)}\}$ at the next resolution level.

- **Trajectory awareness:** Each tile carries a trajectory annotation $\tau_i^{(k)}$: the record of prior state sequences that have passed through the tile's domain.

- **Noise annotation:** Each tile carries a noise annotation $\eta_i^{(k)}$: the record of unresolved structure within the tile that could not be assigned to any trajectory.

The fundamental departure from standard multi-resolution analysis is the noise annotation. Rather than discarding unresolved structure as error to be minimized, TARTAN archives it as annotated residue available for future reinterpretation. This operationalizes the hallucination-as-normal principle of Chapter 2: the unresolved structure is the non-trivial H^1 cocycle of the local sheaf.

13.21 The Annotated Noise Schema

Each noise annotation $\eta_i^{(k)}$ carries the following metadata:

- **Origin:** Which trajectory segment produced this unresolved structure.
- **Constraint violation type:** Which admissibility condition the structure failed (thermodynamic, logical, geometric, historical).
- **Reinterpretation candidates:** The set of alternative trajectory segments that could absorb this structure under

relaxed gluing conditions.

- **Persistence weight:** How long the noise has remained unresolved; high-persistence noise is prioritized for repair.

The annotated noise schema converts gluing failures from silent errors into structured records. A TARTAN system operating over a long time horizon accumulates an annotated noise ledger that functions as a diagnostic tool: recurring noise patterns indicate persistent admissibility failures that require structural repair rather than local patching.

13.22 Local Plausibility versus Global Consistency

The central insight of TARTAN is that local plausibility and global [H] consistency are distinct properties that standard multi-resolution methods conflate. A locally plausible configuration is one that satisfies the admissibility conditions within a single tile. A globally consistent configuration is one whose local solutions glue into a coherent global section.

Definition 13.10 (TARTAN Consistency Check)

A tiling $\mathcal{T}$ is *globally consistent at level* k if for every pair of adjacent tiles $(T_i^{(k)}, T_j^{(k)})$ sharing a boundary ∂_{ij}:

$$\tau_i^{(k)}\big|_{\partial_{ij}} = \tau_j^{(k)}\big|_{\partial_{ij}},$$

that is, the trajectory annotations agree on the shared boundary. A failure of this condition is a *seam inconsistency*, recorded in the noise annotation of both tiles.

13.23 Four Projection Modes

TARTAN operates in four projection modes corresponding to the four derived stacks of the RSVP field system (Appendix 13.5):

Mode 1: Thermodynamic projection. Each tile tracks the local entropy production rate. Seam inconsistencies generate entropy; repairs reduce it. The tiling evolves by preferentially extending tiles that reduce overall entropy production.

Mode 2: Kinetic projection. Each tile tracks the local velocity field. Seam inconsistencies manifest as velocity discontinuities at tile boundaries. The tiling evolves by smoothing velocity across boundaries while preserving trajectory annotation.

Mode 3: Informational projection. Each tile tracks the local residue field $\mathcal{R}$. Seam inconsistencies manifest as conflicting provenance records at tile boundaries. The tiling evolves by resolving provenance conflicts through the sheaf repair operations of Chapter 6.

Mode 4: Semantic projection. Each tile tracks the local semantic density Φ and its annotation. Seam inconsistencies are exactly the semantic tears of the Yarncrawler framework—mismatches between adjacent local interpretations. The tiling evolves by the stigmergic repair mechanism of Theorem 6.4.

13.24 Relationship to Hallucination Formalism

Under TARTAN, hallucination is precisely a seam inconsistency [H] that has been filled by a locally plausible but globally inconsistent trajectory annotation. The annotation satisfies the admissibility conditions within its tile but fails the global consistency check at tile boundaries. The annotated noise schema records this failure with its reinterpretation candidates, enabling future repair when additional evidence arrives.

This makes the corrigibility condition of Definition 2.3 computationally explicit: a corrigible hallucination is one whose seam inconsistency is recorded in the noise annotation with reinterpretation candidates; a drifting hallucination is one whose seam inconsistency has been silently propagated to adjacent tiles without annotation, corrupting their trajectory records.

CLIO: Constraint-Leveraged Inference and Optimization

CLIO— Constraint-Leveraged Inference and Optimization — is the recursive constraint repair framework developed as the self-correction layer of the Yarncrawler architecture. This appendix distinguishes the framework from the convergent system of Cheng, Broadbent, and Chappell tepcheng2025clio, develops the formal specification, and identifies the key design differences.

13.25 Two Versions of CLIO

The acronym CLIO has been used independently in two frameworks with overlapping but distinct concerns. The Cheng–Broadbent–Chap system (hereafter CBC-CLIO) is an inference-time orchestration layer over large language models that uses recursive confidence signals to steer reasoning without post-training. The framework developed here (hereafter RC-CLIO for Recursive Constraint repair) addresses a different problem: maintaining coherence across overlapping theoretical frameworks in a distributed conceptual manifold.

The convergences are real. Both systems use recursive depth as a control variable. Both treat uncertainty reduction as a positive signal. Both expose their reasoning process rather than hiding it behind post-hoc rationalization. The CBC-CLIO finding that correct reasoning exhibits a negative uncertainty gradient over time is consistent with the RC-CLIO corrigibility condition: a system whose reconstruction errors are being corrected shows decreasing unresolved residue.

The divergences are equally real. CBC-CLIO optimizes toward known correct answers on benchmarks. RC-CLIO operates in settings where the admissibility conditions themselves are being constructed rather than given—where there is no external oracle to verify the output. CBC-CLIO's uncertainty signal is self-reported by the model. RC-CLIO's uncertainty signal is the non-trivial cohomological residue of the theoretical manifold, detectable through gluing failures at framework boundaries.

13.26 The RC-CLIO Specification

Definition 13.11 (RC-CLIO Repair Cycle)

A single *RC-CLIO repair cycle* over a theoretical manifold $(\mathcal{M}, \{(U_i, \varphi_i)\})$ consists of four phases:

Phase 1: Seam detection. Compute the seam penalty $L_{\text{seam}}(x) = \sum_{i<j} w_i(x) w_j(x) \|f_i(x) - f_j(x)\|^2$ over all overlapping charts. Identify regions where L_{seam} exceeds a

threshold ϵ_{repair}.

Phase 2: Cocycle identification. For each high-seam region, identify the Čech 1-cocycle responsible: the mismatch between restrictions of adjacent local sections to their shared overlap.

Phase 3: Repair morphism construction. Introduce new morphisms into the affected local sections that trivialize the cocycle. The repair morphism must satisfy locality (modifying only sections near the tear), type safety (remaining within the admissible semantic category), and stigmergic monotonicity (reducing the seam penalty without increasing it elsewhere).

Phase 4: Residue annotation. Record the repair in the system's provenance structure: which seam was repaired, by which morphism, and what alternative repairs were available but not applied.

13.27 The Corrigibility Condition in RC-CLIO

Definition 13.12 (RC-CLIO Corrigibility)

An RC-CLIO system is *corrigible* with respect to a target manifold $\mathcal{M}^*$ if:

1. The seam detection phase correctly identifies all regions where the current manifold $\mathcal{M}$ differs from $\mathcal{M}^*$.
2. The repair morphisms introduced in Phase 3 reduce $d(\mathcal{M}, \mathcal{M}^*)$ monotonically.
3. The residue annotation in Phase 4 is complete: all applied morphisms are recorded with their alternatives.

A system fails corrigibility if it introduces repair morphisms that reduce the seam penalty L_{seam} while increasing $d(\mathcal{M}, \mathcal{M}^*)$: this is proxy stabilization at the level of the repair process itself.

13.28 Relationship to CBC-CLIO

The CBC-CLIO system implements a version of Phase 1 (seam detection through uncertainty monitoring) and Phase 3 (repair through recursive re-invocation of the reasoning process). It does not implement Phase 2 (cohomological identification of the obstruction) or Phase 4 (provenance annotation of repairs). This is appropriate given CBC-CLIO's operational context: it is designed for single-session scientific question-answering, not for long-horizon theoretical manifold maintenance.

[C] The CBC-CLIO uncertainty oscillation finding—that correct answers show negative uncertainty gradients while incorrect answers show positive or oscillating gradients—is consistent with Phase 1 of RC-CLIO. In RC-CLIO terms: a system whose seam penalties are converging to zero is on a corrigible trajectory; a system whose seam penalties are oscillating has entered a repair cycle that is not reducing $d(\mathcal{M}, \mathcal{M}^*)$ despite reducing L_{seam} locally. The oscillation is the signature of proxy stabilization in the repair process.

Obstruction Cohomology and Katz-Admissibility

This appendix develops the mathematical background for the sheaf-theoretic institutional failure analysis of Chapter 9. The central question is: under what conditions do locally consistent solutions resist extension to globally consistent ones? The answer is given by obstruction cohomology.

13.29 The Local-to-Global Extension Problem

Let X be a topological space modeling an institutional domain (a regulatory jurisdiction, a multi-department organization, or a federated system of governance). Let $\mathcal{U} = \{U_i\}_{i \in I}$ be a cover of X by overlapping local domains.

Definition 13.13 (Local Solution)

A *local solution* on U_i is a section $s_i \in \mathcal{S}(U_i)$ of a sheaf $\mathcal{S}$ that satisfies the operational requirements specified for domain

U_i. A collection of local solutions $\{s_i\}_{i \in I}$ is *compatible* on overlaps if for all i, j:

$$\rho^{U_i}_{U_i \cap U_j}(s_i) = \rho^{U_j}_{U_i \cap U_j}(s_j),$$

where ρ denotes the restriction maps of the sheaf.

A globally consistent solution is a global section $s \in \mathcal{S}(X)$ restricting to each s_i. The central theorem is that compatible local solutions do not always extend to global sections: the obstruction is measured by the first cohomology group $H^1(X, \mathcal{S})$.

13.30 The Čech Cohomology Construction

For a cover $\mathcal{U}$, define the Čech cochain groups:

$$C^0(\mathcal{U}, \mathcal{S}) = \prod_i \mathcal{S}(U_i), \tag{13.5}$$

$$C^1(\mathcal{U}, \mathcal{S}) = \prod_{i<j} \mathcal{S}(U_i \cap U_j), \tag{13.6}$$

$$C^2(\mathcal{U}, \mathcal{S}) = \prod_{i<j<k} \mathcal{S}(U_i \cap U_j \cap U_k). \tag{13.7}$$

The coboundary map $\delta : C^0 \to C^1$ is defined by:

$$(\delta s)_{ij} = \rho^{U_j}_{U_i \cap U_j}(s_j) - \rho^{U_i}_{U_i \cap U_j}(s_i).$$

A collection of local solutions is compatible if and only if $\delta s = 0$, that is, it is a Čech 0-cocycle. The global extension exists if and only if the cocycle is also a coboundary: $s = \delta t$ for some $t \in C^{-1}$ (which reduces to the condition that the collection arises from a global section).

The obstruction group is:

$$H^1(\mathcal{U}, \mathcal{S}) = \ker(\delta : C^1 \to C^2)/\mathrm{im}(\delta : C^0 \to C^1).$$

A non-trivial element of H^1 is a compatible family of local solutions that cannot be assembled into a global section. It is exactly a Katz-inadmissible configuration.

13.31 Katz-Admissibility Revisited

Definition 13.14 (Katz-Admissibility (Formal))

An institutional configuration $\{s_i\}_{i \in I}$ is *Katz-admissible* if its associated Čech 1-cocycle $[\delta s] \in H^1(\mathcal{U}, \mathcal{S})$ is trivial. Equivalently, the local solutions extend to a global section of $\mathcal{S}$ over the entire institutional domain X.

An institution is *Katz-inadmissible* if $H^1(\mathcal{U}, \mathcal{S})$ is non-trivial and the current configuration represents a non-trivial element of this group.

> *Proposition 13.2* (*Persistence of Katz-Inadmissibility*)
>
> If an institutional configuration is Katz-inadmissible, then no local reform within any single domain U_i can resolve the obstruction. Resolution requires either: (i) modifying the sheaf $\mathcal{S}$ (changing the compatibility requirements themselves), (ii) refining the cover $\mathcal{U}$ (restructuring domain boundaries), or (iii) accepting that no globally consistent operation exists and managing the resulting incoherence explicitly.

Sketch. A non-trivial element of H^1 is by definition not in the image of $\delta : C^0 \to C^1$. Modifying s_i within U_i alone changes the C^0 element but does not change the cohomology class of the resulting 1-cocycle, since the class depends on the mismatch at overlaps, not on the individual sections. Local reform within U_i can change s_i but cannot change the compatibility structure at $U_i \cap U_j$ if s_j is not simultaneously reformed. And if s_j must also be reformed, the problem has escaped the domain of local reform. $\square$

13.32 Higher Obstructions

[C] The higher cohomology groups H^k for $k \geq 2$ encode deeper structural obstructions. A non-trivial H^2 class indicates that even after resolving all pairwise overlaps between domains, the triple overlaps $U_i \cap U_j \cap U_k$ impose additional compatibility requirements that resist satisfaction simultaneously.

In institutional terms, H^2 obstructions correspond to cases where three departments can each be pairwise reconciled but cannot all three operate consistently simultaneously—coordination failures that appear only when three-way interaction is considered. Such obstructions can exist even when every pairwise conflict has been locally resolved.

13.33 Worked Example: A Three-Department Institution

Let X be an institution with three departments A, B, C, each covering a domain U_A, U_B, U_C with overlaps. Suppose: [P]

- A and B are compatible on $U_A \cap U_B$: their operational procedures agree.
- B and C are compatible on $U_B \cap U_C$.
- A and C are compatible on $U_A \cap U_C$.

This appears to be a situation with no conflicts. Yet if the sheaf $\mathcal{S}$ is non-abelian (if the order of operations matters), pairwise compatibility does not imply global consistency. The triple overlap $U_A \cap U_B \cap U_C$ may impose a consistency condition that the pairwise agreements jointly violate.

The resulting H^2 obstruction explains why the institution experiences recurring coordination failures that no bilateral negotiation can resolve: the problem is topological, not interpersonal.

Juarrero and Deacon: Framework Comparisons

This appendix situates the process ontology of this monograph relative to two major contemporary accounts of emergence, constraint causation, and intentionality: Alicia Juarrero's constraint-based dynamics and Terrence Deacon's absential organization. Both influence the framework developed here; neither is adopted wholesale. The differences are philosophical load-bearing rather than terminological.

13.34 Juarrero: Context-Sensitive Constraints

Juarrero's central contribution in *Dynamics in Action* (Juarrero, 1999) is the distinction between first-order and second-order constraints. First-order constraints are boundary conditions that limit the range of possible system states directly. Second-order constraints are context-sensitive: they alter the probability distribution over future states rather than eliminating possibilities outright. Juarrero argues that intentional behavior is best understood as second-order constraint causation: the meaning of an action is determined by the context that constrains which actions are admissible, not by the intrinsic properties of the action itself.

Points of convergence with the present framework:

The trajectory-admissibility framework is continuous with Juarrero's account. Admissible trajectories are exactly the states that survive the imposition of both first-order and second-order constraints. The historical accumulation of constraints—the decreasing filtration $\mathcal{A}(\mathcal{C}_0) \supseteq \mathcal{A}(\mathcal{C}_1) \supseteq \dots$—is a formalization of what Juarrero means by context-sensitive constraint accumulation. The RSVP entropy field S functions as a measure of residual second-order constraint: regions of high S are regions where the contextual constraints have not yet settled into a stable pattern.

Points of divergence:

Juarrero's account sometimes approaches a strong reading on which second-order constraints are themselves causally productive—on which the context does not merely filter but actively contributes causal force to the outcomes it shapes. The present framework takes a more deflationary position: constraints are eliminative rather than productive. They remove inadmissible trajectories but do not add new causal forces. Apparent goal-directedness is a consequence of what has been eliminated, not of any positive force directed toward a target. This is a subtle but important difference in the direction of causation.

13.35 Deacon: Absential Organization

Deacon's account in *Incomplete Nature* (Deacon, 2011) introduces the concept of *absence* as a causally efficacious feature of complex systems. The central claim is that teleodynamic organization—the kind that characterizes living systems and minds—is constituted by the causal influence of what is not present: potential states toward which the system is organized, constraints that do not currently exist but that the system's organization tends to produce.

Deacon develops a hierarchy of emergent dynamics:

- **Homeodynamics**: simple attractor dynamics in physical systems.
- **Morphodynamics**: self-organizing systems that develop global regularities from local interactions.
- **Teleodynamics**: systems whose organization is constituted by reference to absent or potential states.

Points of convergence with the present framework:

The teleodynamics level is genuinely important for understanding the kind of coherence exhibited by living systems and minds. Deacon is right that adequate accounts of intentionality require more than simple homeodynamics or morphodynamics. The KES map's successor hypothesis space H_{t+1} is exactly an absent state that constrains present dynamics: the system's current trajectory is partly constituted by which future hypothesis spaces remain admissible.

Points of divergence:

The present framework declines to grant absence ontological primacy. On Deacon's strongest reading, not-yet-realized states exert genuine causal influence on current dynamics—there is a kind of backward causation or final causation operating through the organization of living systems. The present framework explains the same phenomena without this commitment. The KES successor hypothesis space H_{t+1} constrains current dynamics not because it is causally efficacious as an absence but because the current constraint set $\mathcal{C}_t$ has already eliminated the trajectories that would make H_{t+1} unreachable. The direction of causation is entirely forward: prior constraints eliminate future possibilities; the future does not pull the present.

Formally: the filtration $\mathcal{A}(\mathcal{C}_0) \supseteq \mathcal{A}(\mathcal{C}_1) \supseteq \ldots$ produces apparent goal-directedness through elimination alone. No backward causation is required. The system appears to be directed toward the configurations that survive the filtration because all alternatives have been progressively eliminated. This is sufficient to generate the phenomena Deacon attributes to teleodynamics without requiring absence to be causally efficacious.

13.36 The Constraint-Final Causation Distinction

The formal distinction between constraint causation and final causation is:

Definition 13.15 (Constraint Causation)

A system exhibits *constraint causation* if its directional structure is fully explained by the accumulation of prior constraints that eliminate inadmissible trajectories. The causal influence flows entirely forward: past constraint events eliminate future possibilities, producing a monotonically narrowing admissibility manifold.

Definition 13.16 (Final Causation)

A system exhibits *final causation* if a not-yet-realized future state s^* exerts causal influence on current system dynamics, such that the dynamics would be different if s^* were different, holding all prior history constant.

The present framework adopts constraint causation throughout and declines final causation. This position is consistent with [C] thermodynamics, bounded computation, and path dependence. It produces all the observed phenomena (apparent purposiveness, anticipatory behavior, structural persistence) without positing causation that runs backward through time or that requires future states to exist ontologically before they are realized.

13.37 Where Deacon's Framework Remains Valuable

Despite the divergence on final causation, Deacon's hierarchical emergence framework provides useful vocabulary for distinguishing levels of organization. The homeodynamic/morphodynamic/teleodynamic hierarchy maps onto the trajectory framework as follows:

Homeodynamic systems navigate a fixed admissibility manifold whose topology is specified in advance. Morphodynamic systems modify their admissibility manifold through self-organization, but the modification rules are themselves fixed. Teleodynamic systems—and this is the genuinely novel level—modify both their admissibility manifold and their modification rules through ongoing interaction with the outcomes of their prior modifications.

The Yarncrawler framework captures exactly this last level: the self-rewrite endofunctor $\mathsf{R} : \mathcal{M} \to \mathcal{M}$ modifies the manifold, while the RC-CLIO repair cycle modifies the repair rules themselves. This is teleodynamic organization in Deacon's sense, achieved through forward-only constraint causation rather than through final causation.

The Stone Piano Hypothesis: An Archaeological Case Study

Whatever they were doing, they intended it.

The Stone Piano Hypothesis proposes that the Bruniquel Cave stalagmite structures, dated to approximately 176,500 years before present tepjaubert2016 and attributed to *Homo neanderthalensis*, functioned as lithophones: percussion instruments fabricated from resonant calcite speleothems. This appendix situates the hypothesis within the Yarncrawler framework.

13.38 Admissibility Filtering in Speleothem Selection

The fundamental frequency of a free-free calcite rod of length L and radius r:

$$f1 = \frac{\lambda1^2\, r}{8\pi L^2}\sqrt{\frac{E}{\rho}}, \qquad \lambda1 \approx 4.730,$$

with $E \approx 75\,\mathrm{GPa}$ and $\rho \approx 2600\,\mathrm{kg\,m^{-3}}$. Fracturing stalagmites from their bases converts clamped-free resonators (inharmonic series $1 : 6.27 : 17.55$) into free-free resonators with more harmonic profiles $(1 : 2.76 : 5.40)$. This is an admissibility filtering operation.

13.39 Sympathetic Resonance as the Corrigibility Mechanism

Acoustic beating—the periodic amplitude modulation produced by two nearly coincident frequencies—disappears when frequencies are matched. Its disappearance is an unambiguous perceptual event providing real-time corrigibility feedback. The builders could tune fragment lengths by progressive shortening, without abstract acoustic knowledge, using only the disappearance of beating as a signal. This is the KES loop operating in stone at 176,500 BP.

13.40 The Ring as Stigmergic Structure

A partially assembled ring of resonant elements constitutes an acoustic admissibility filter. New candidate fragments that excite strong sympathetic vibrations in the existing ring are harmonically admissible; those that do not are refused. The structure under construction participates in selecting the components from which it is built: Axiom 3 (Autocatalytic Fix Design) instantiated in deep prehistory.

13.41 Estimated Chamber Reverberation

By Sabine's formula with $V \approx 73\,\mathrm{m}^3$ and $\alpha \approx 0.03$:

$$T60 \approx \frac{0.161 \times 73}{0.03 \times 130} \approx 3.0\,\mathrm{s}.$$

The cave amplifies and extends each repair operation, lowering the threshold for subsequent ones. It is a Yarncrawler environment.

Companion Document Index

The following companion documents are cross-referenced throughout this monograph. All are available at `https://github.com/standa`

Physics and Cosmology: *Axioms for a Falling Universe*; *The Persistence of Structure*; *Admissible Geometry*; *Fossil Manifolds and Mediated Extraction*; *Geometry from Entropy Flow*; *Physics After Spacetime*; *Scalar Irruption via Entropic Differential*.

Cognition and Epistemology: *The Collapse of Epistemic Coherence*; *Hallucination is Normal*; *Never Predict Noise*; *World-State Reconstruction*; *Constraint Before Completion*; *Selection Without Design*; *The Compiled Self*; *Noun-Free Cognition*.

Computation: *Yarncrawler in Action*; *The Elements of Computational Worlds*; *Geometry, Cognition, and the Transparency of Computation*; *Semantic Relaxation Networks*; *Trajectories All the Way Down*; *Irreversibility as Architecture*.

Social and Institutional: *The Collapse of Proxy Integrity*; *Participation without Guarantee*; *Recomposable Fragmentation*; *The Entropy of Austerity*; *Obstruction Cohomology*; *Against Namespace Laundering*; *The Stack Capture Race*.

Paleoarchaeology: *The Stone Piano Hypothesis*.

Perceptual Reconstruction and Harmonic Completion

> A signal may remain semantically recognizable while losing fine-grained structural relations necessary for perceptual richness, temporal immersion, or sustained attentional coherence.

Perception does not operate through perfect recovery of external structure. All reconstructive systems operate under informational incompleteness, energetic limitation, and finite temporal accessibility. Perceptual systems therefore rely upon partial reconstruction procedures that stabilize coherent interpretations from degraded or underspecified inputs. This appendix investigates harmonic reconstruction, perceptual coherence, and temporal continuity under compressed representation regimes.

The material here should be understood as **[P]** phenomenological interpretation and **[H]** heuristic engineering throughout. It does not establish that compressed audio damages consciousness, removes emotional essence, or eliminates hidden biological frequencies necessary for cognition. Such claims exceed available evidence. The framework proposes only that reconstruction quality influences perceptual coherence and that harmonic continuity may affect subjective richness, immersion, and attentional stability under constrained representational conditions.

13.42 Compression and the Symbolic/Inferential Distinction

All lossy compression systems preserve certain structural relations while discarding others. The resulting reconstruction is shaped not only by retained information but also by the assumptions embedded within the reconstruction procedure itself.

Perceptual degradation frequently emerges when temporally or harmonically significant relations are removed despite preserving coarse symbolic intelligibility. This reflects the broader separation between symbolic continuity and inferential continuity developed in Chapter 3. A signal may remain semantically recognizable—the words are audible, the melody is identifiable—while losing fine-grained structural relations necessary for perceptual richness, temporal immersion, or sustained attentional coherence.

13.43 Harmonic Reconstruction as Constraint Completion

Suppose a signal admits a partial harmonic decomposition:

$$x(t) = \sum_{k=1}^{N} a_k \cos(\omega_k t + \phi_k).$$

Compression removes a subset of components:

$$\hat{x}(t) = \sum_{k \in S} a_k \cos(\omega_k t + \phi_k), \quad S \subseteq \{1, \ldots, N\}.$$

Reconstruction seeks admissible completion operators $\mathcal{R} : \hat{x}(t) \mapsto \tilde{x}(t)$ that infer missing relational structure from surviving harmonic constraints:

$$\tilde{x} = \arg\min_{y \in \mathcal{C}} D(y, \hat{x}),$$

where $\mathcal{C}$ denotes the space of temporally coherent reconstructions and D extracts perceptually relevant relational structure rather than raw amplitude similarity alone.

The perceptual divergence between original and reconstruction is:

$$D(x, \tilde{x}) = \int \|K(x(t)) - K(\tilde{x}(t))\|^2 \, dt,$$

where K extracts perceptually relevant relational structure. This is a constrained interpolation problem, not a restoration problem. The reconstructed signal is an inferential approximation generated under bounded informational accessibility, not a recovery of the original.

13.44 Temporal Continuity and Perceptual Stability

Biological perception exhibits strong sensitivity to temporal discontinuity. Abrupt fragmentation, phase instability, and incoherent transient structure increase perceptual fatigue and reduce sustained attentional stabilization. Define local temporal fragmentation through phase structure $\phi(t)$:

$$F(t) = \left| \frac{d^2\phi}{dt^2} \right|.$$

Perceptual stability emerges when integrated fragmentation remains bounded:

$$\int_{t_0}^{t_1} F(t)\,dt < \epsilon.$$

Harmonic reconstruction mechanisms may partially restore such continuity even when substantial information has been lost. The resulting increase in perceived richness does not imply that hidden information has been recovered in a literal sense. Rather, the reconstruction process stabilizes relational structures [H] compatible with prior perceptual expectations and surviving signal constraints.

13.45 Inferential Completion in Auditory Perception

Perceptual reconstruction depends heavily upon inferential completion. Human auditory systems routinely interpolate missing information from contextual structure, harmonic expectation, and temporal continuity:

$$\hat{x} = \arg\max_{z} P(z \mid y, H),$$

where $y \subset x$ is the available perceptual input and H represents prior structural expectations. Hallucination, interpolation, and perceptual completion all arise naturally whenever $\dim(y) < \dim(x)$. The system operates by constrained reconstruction rather than exhaustive recovery of inaccessible total states. This is exactly the reconstruction bound of Chapter 2 applied to the auditory domain.

[P] Such effects belong primarily to psychoacoustics, perceptual inference, and temporal coherence theory. They do not require speculative metaphysics about consciousness or hidden frequencies. The interesting questions are computational and psychoacoustic: which structural relations does the auditory system most depend on for stability, and which reconstruction heuristics best preserve those relations under compression?

Bibliography

Deacon, T. W. (2011). *Incomplete Nature: How Mind Emerged from Matter*. W.W. Norton, New York.

Flyxion (2026a). Axioms for a falling universe: A unified RSVP lagrangian. In *Companion Documents*. Available at `https://standardgalactic.github.io/alphabet/projects/axioms-for-a-falling-universe.pdf`.

Flyxion (2026b). Yarncrawler in action. In *Companion Documents*. Available at `https://standardgalactic.github.io/yarncrawler/Yarncrawler%20in%20Action.pdf`.

Juarrero, A. (1999). *Dynamics in Action: Intentional Behavior as a Complex System*. MIT Press, Cambridge, MA.