The Admissibility Principle

Computation, Physics, and Analogy as One Operation

Flyxion

Independent Researcher

May 2026

Abstract

We propose that a single abstract operation underlies β-reduction in the λ-calculus, membrane collapse in the Spherepop computational geometry, analogy-driven chunking in Hofstadter's theory of cognition, cellular differentiation as reconstructed through multimodal lineage tracing (Wang, He, & Hu 2026), entropy-constrained field relaxation in the Relativistic Scalar-Vector Plenum (RSVP) framework, quantum interference-mediated dwell-time anomalies, phase-matched energy redistribution in ultrafast nonlinear optics, maximum-dissipation selection in turbulent fluid dynamics, recursive epistemic admissibility in strategic games, and temporal-domain feature extraction in the Marine signal architecture. We call this operation *admissibility-constrained collapse*: a bounded region accepts a compatible perturbation, internally reorganizes its constraint structure, and collapses into a new admissible successor state while preserving certain invariants. The paper formalizes a generalized reduction operator $\mathcal{R}$ of which each domain provides a specialization, and identifies a family of local-compatibility reconciliation structures — exemplified by Church–Rosser confluence, sheaf gluing, and iterated admissibility elimination — that recur across domains as instances of a shared convergence grammar rather than a single proved universal theorem. Spherepop is situated as a proposed geometric operational semantics mediating between the symbolic (λ-calculus), cognitive (Hofstadter), and thermodynamic (RSVP) levels, and its current status as a research program awaiting axiomatization is acknowledged explicitly. The biological case of cellular differentiation is identified as the strongest empirical pillar: lineage reconstruction requires historical admissibility propagation as a computational necessity, not as a theoretical overlay. Epistemic status is stratified throughout into five explicit layers: established peer-reviewed results, engineering prototypes, formal abstraction, philosophical synthesis, and open research program.

Contents

1. Introduction and Motivation

1.1. The central claim

What does a photon do when it briefly becomes an atomic excitation? What does a λ-term do when it swallows its argument? What does a mind do when a new perception activates a stored conceptual chunk? What does a nonlinear crystal do when a chirped pulse traverses it? What does a pluripotent stem cell do when it commits to a developmental lineage?

Superficially these are questions from different sciences with different vocabularies. We argue they share a single formal skeleton: a bounded region equipped with a constraint structure accepts an incoming perturbation, tests its local compatibility, and collapses into a new configuration while preserving an invariant. We call this the *admissibility principle*.

The principle has at least four realizations that can be made formally precise and at least three more that we treat as motivated engineering architectures or speculative proposals. These realizations are not physically identical, and the paper is careful not to conflate different evidential strata. The unity we claim is structural, not material: the same abstract grammar $\mathcal{R}(\mathcal{B}_i, \Delta_i) \to \mathcal{B}_{i+1}$ governs each case at an appropriate level of abstraction.

1.2. Why this matters

Standard cross-disciplinary syntheses suffer one of two failure modes. The first is *premature literalism*: claiming that disparate systems are physically the same or that metaphors constitute proofs. The second is *vacuous analogy*: the observation that "everything transforms" has no predictive content.

The admissibility principle attempts to occupy the productive middle ground by requiring, for each domain, an explicit identification of:

- the admissibility region $\mathcal{B}_i$ and its constraint structure;
- the perturbation Δ_i and its compatibility condition;
- the collapse operator $\mathcal{R}$ and what it preserves;
- the invariant or normal form $\mathcal{B}_{i+1}$.

Where these four elements can be made precise, the analogy is a theorem or a well-defined conjecture. Where they cannot, the analogy remains heuristic and is labeled accordingly.

1.3. What this framework does not claim

It is important to state explicitly what the admissibility principle is *not*.

- It does not claim that cognition is literally λ-calculus, or that turbulence is secretly semantics, or that RSVP parameters are established physical constants.
- It does not claim that all admissibility systems are mathematically identical or that the generalized reduction operator $\mathcal{R}$ is a proven theorem rather than a research schema.
- It does not claim that every constrained process instantiates $\mathcal{R}$. A thermostat constrains temperature but does not qualify: it filters outputs without recursively restructuring the state space. A lookup table maps inputs to outputs without propagating a transportable invariant. Mere constraint does not entail admissibility-collapse.
- It does not claim formal equivalence between domains that are only structurally analogous. The paper distinguishes throughout: formal equivalence (provable), structural analogy (argued), admissibility-preserving transport (conjectured), and metaphorical correspondence (labeled as such).

What the framework *does* claim is that many systems exhibit recurring structural motifs involving constrained trajectory reduction, invariant preservation, and local compatibility propagation, and that these motifs are sufficiently precise to motivate a shared research grammar. The analogy to category theory is apt: category theory did not claim all mathematics was literally one object, but supplied a transport language for structural relationships across domains. The admissibility principle aims to play the same role for constraint-preserving collapse.

1.4. Organization

The paper is organized into four parts plus a discussion, followed by seven appendices.

Part I (Sections 2–5) develops the formally precise core: λ-calculus and the Curry–Howard correspondence, Spherepop as proposed geometric operational semantics, Hofstadter's theory of analogy as the cognitive reading, and the formulation of analogy itself as a reduction operation.

Part II (Section 6) provides the biological instantiation: cellular differentiation reconstructed through multimodal lineage tracing, drawn from Wang, He, and Hu (2026) in *Nature Reviews Genetics.* This is the domain in which admissibility is a computational necessity rather than an interpretive overlay, and it serves as the strongest empirical pillar of the monograph.

Part III (Sections 7–8) treats established physics: quantum weak-value dwell times and DC-OPA ultrafast optics, each mapped onto the admissibility schema.

Part IV (Sections 9–10) develops the speculative and prototypical layers: the RSVP thermodynamic framework, the Marine/Phoenix 8a signal architecture, and the MEM|8 semantic memory proposal, all explicitly marked as conjectural or engineering-stage.

Part V (Sections 11–16) synthesizes all layers into the generalized reduction operator, develops the game-theoretic, turbulence, and relativistic admissibility cases, articulates the five-layer epistemic architecture, and integrates RSVP, CLIO, TARTAN, and Yarncrawler.

Section 17 identifies open problems, distinguishes what has been established from what remains to be done, and states the framework's self-assessment as a research grammar rather than a unified theory. Section 18 develops the underconstraint problem following Yudkowsky (2008) and connects it to the sheaf-theoretic obstruction framework. The seven appendices develop the category-theoretic, fiber-bundle, sheaf-cohomological, variational, homotopy-theoretic, non-vacuity, and three-level coarse-graining formulations in mathematical detail.

Part I

The Formal Core

2. λ-Calculus, β-Reduction, and Curry–Howard

Epistemic Status: Established (peer-reviewed results)

The results in this section are standard and fully established. Primary sources: Church (1936), Curry and Feys (1958), Howard (1980), Girard (1986), and Sørensen–Urzyczyn (2006).

2.1. Syntax and reduction

The untyped λ-calculus is defined over the grammar

$$M, N \ ::= \ x \mid \lambda x.\, M \mid (M\, N),$$

where x ranges over a countably infinite set of variables. The canonical computation rule is *β-reduction*:

$$((\lambda x.\, M)\, N) \ \to_\beta \ M[x := N]. \tag{1}$$

The substitution $M[x := N]$ replaces every free occurrence of x in M with N, with the standard capture-avoidance proviso (renaming bound variables, α-conversion, as necessary).

Definition 2.1 (Admissibility in λ-calculus)**.** A redex $((\lambda x.\, M)\, N)$ is *admissible* if N is free for x in M modulo α-conversion. Every redex is admissible after appropriate renaming.

The reduction relation $\to_\beta^*$ is the reflexive transitive closure of $\to_\beta$. A term M is in *β-normal form* if it contains no redex.

2.2. Church–Rosser confluence

Theorem 2.2 (Church–Rosser, 1936). *If $M \to_\beta^* M_1$ and $M \to_\beta^* M_2$, then there exists M_3 such that $M_1 \to_\beta^* M_3$ and $M_2 \to_\beta^* M_3$.*

Theorem 2.2 asserts that distinct reduction paths from a common source can always be *joined*: there exists a common reduct downstream. Normal forms, when they exist, are therefore unique up to α-equivalence. This is the first instance of the master convergence principle that we will trace across all domains.

2.3. Curry–Howard: types as propositions

The *Curry–Howard correspondence* identifies:

$$\begin{aligned} \text{propositions} &\leftrightarrow \text{types}, \\ \text{proofs} &\leftrightarrow \text{typed terms}, \\ \text{proof normalization} &\leftrightarrow \beta\text{-reduction}, \\ \text{provability} &\leftrightarrow \text{type inhabitation}. \end{aligned}$$

Under the simply-typed λ-calculus, every type is a proposition of minimal implicational logic and every well-typed term is a proof. Normalization of proofs corresponds exactly to β-reduction of terms.

Remark 2.3. The BHK (Brouwer–Heyting–Kolmogorov) interpretation provides the semantic backing: a proof of $A \Rightarrow B$ is a *procedure* that transforms proofs of A into proofs of B. This already anticipates the admissibility principle: a type/proposition is a constraint, a term/proof is the compatible perturbation, and normalization is the collapse into normal form.

2.4. Divergence and fixed points

In the untyped setting, not all terms normalize. The combinator

$$\Omega = ((\lambda x.\,(x\,x))\,(\lambda x.\,(x\,x)))$$

diverges: $\Omega \to_\beta \Omega \to_\beta \cdots$. The fixed-point combinator

$$Y = \lambda f.\,((\lambda x.\,(f\,(x\,x)))\,(\lambda x.\,(f\,(x\,x))))$$

satisfies $(F\,((Y\,F))) =_\beta (Y\,F)$ for any F. Divergence in our schema corresponds to an admissibility loop with no stable attractor: the system perpetually reorganizes without collapsing to a normal form.

3. Spherepop: Geometric Operational Semantics

Epistemic Status: Speculative Theoretical Proposal

Spherepop is a theoretical computational framework under development. The correspondence with λ-calculus stated here is a formal proposal, not an independently verified result. The claim is that the mapping in Table 1 constitutes a sound interpretation; proving this requires a complete formal specification of Spherepop syntax and reduction rules, which is ongoing work.

3.1. Motivation

Standard λ-calculus is syntactic: computation is symbol manipulation governed by string-rewriting rules. Substitution $M[x := N]$ is a meta-theoretic operation, not an object-level entity.

Spherepop proposes to *externalize* substitution as a spatial, topological, and historical operation. The guiding intuition: a λ-abstraction creates a *protected evaluation region*; application injects an argument into that region; substitution propagates through it; the region boundary dissolves once computation is resolved. This sequence is not a metaphor — it is a proposed formal operational semantics in which boundaries are first-class objects.

3.2. Core vocabulary

Definition 3.1 (Bubble)**.** A *bubble* $\mathbf{b} = (I, \partial, H)$ consists of:

- an *interior* I: a set of latent sub-expressions;
- a *membrane* ∂: a constraint structure specifying what perturbations are locally admissible;
- a *history* H: an ordered record of prior reductions that produced the current state.

Definition 3.2 (Admissible injection)**.** A bubble $\mathbf{b}'$ is *admissible* for injection into $\mathbf{b}$ if $\mathbf{b}'$ is compatible with $\partial(\mathbf{b})$ in the sense that all free attachment sites in $I(\mathbf{b})$ can be bound by the structure of $\mathbf{b}'$ without capturing existing bindings in $H(\mathbf{b})$.

Definition 3.3 (Membrane collapse)**.** Given an admissible injection of $\mathbf{b}'$ into $\mathbf{b}$, *membrane collapse* is the operation:

$$\mathbf{b} \oplus \mathbf{b}' \;\longrightarrow\; \mathbf{b}''$$

where $\mathbf{b}''$ has:

- interior $I(\mathbf{b}'') = I(\mathbf{b})[\text{open sites} := \mathbf{b}']$,

- history $H(\mathbf{b}'') = H(\mathbf{b}) \cdot \langle \mathbf{b}' \rangle$ (history extended with the injection event),
- membrane $\partial(\mathbf{b}'')$ determined by the residual open sites of $I(\mathbf{b}'')$.

The key distinction from standard substitution: H makes collapse *irreversible* and *identity-preserving.* Two bubbles with identical interiors but different histories are distinct computational objects.

3.3. Correspondence with λ-calculus

Proposition 3.4 (Spherepop–λ correspondence)**.** *Under the interpretation of Table 1, membrane collapse is the geometric realization of β-reduction. α-conversion corresponds to historical rebinding of attachment sites; confluence corresponds to path-independent collapse producing the same stable topology.*

Table 1: Correspondence between λ-calculus and Spherepop.

λ-calculus	**Spherepop**
Variable	Open attachment site
Abstraction	Scoped bubble with membrane
Application	Bubble injection
β-reduction	Membrane collapse
α-conversion	Historical rebinding
β-normal form	Stable bubble configuration
Divergence (Ω)	Infinite recursive bubbling
Confluence (Church–Rosser)	Path-independent topological collapse

3.4. Spherepop as a research program rather than a completed formalism

The reviewer of any draft of this monograph will correctly observe that Spherepop currently lacks a complete operational semantics: there are no published reduction rules, no formal grammar with a proved decidable word problem, no normalization theorem. This observation is accurate and should be acknowledged directly.

Spherepop currently occupies the historical position that manifold theory occupied before rigorous atlases and transition maps were axiomatized, or that tensor notation occupied before Ricci calculus was stabilized. Long before these formalisms were complete, geometers and physicists reasoned productively through intuition and local patching. Spherepop is at that pre-axiomatic stage.

What the present work contributes is the identification of the minimal ingredients that a completed Spherepop formalism would require:

1. *Scoped regions* with formally defined interiors and membranes;
2. *Historical identity persistence* distinguishing bubbles with identical interiors but different causal ancestry;

3. *Admissible injection*, defined by a compatibility predicate on membranes;
4. *Collapse propagation*, specifying how substitution propagates through the interior after injection;
5. *Topology-sensitive rebinding*, the analogue of α-conversion in the spatial setting.

These ingredients correspond exactly to the components of explicit substitution calculi ($\lambda\sigma$, $\lambda\upsilon$), with the addition of irreversible historical extension. Formalizing Spherepop is therefore a well-posed research problem, not an indefinitely open aspiration.

The present paper treats Spherepop as a proposed geometric semantics awaiting axiomatization, not as a completed system. Claims involving Spherepop are accordingly marked *Epistemic Status: Speculative Theoretical Proposal* throughout.

3.5. Explicit substitution and visible dependency

Traditional λ-calculus hides substitution as a meta-operation. Explicit substitution calculi (de Bruijn indices, $\lambda\sigma$, $\lambda\upsilon$) bring it into the object language. Spherepop can be understood as a *spatial explicit substitution calculus*: every bubble carries a visible dependency graph describing what may enter, what is bound, and which historical trajectories remain admissible after collapse. This makes Spherepop closer in spirit to $\lambda\upsilon$ or λx calculi than to the original Church system, but with the additional dimension of irreversible event history.

4. Hofstadter: Analogy as Admissibility-Driven Collapse

Epistemic Status: Established (peer-reviewed results)

Hofstadter's theoretical claims about analogy-making are from published work (*Gödel, Escher, Bach*, 1979; *Fluid Concepts and Creative Analogies*, 1995; *Surfaces and Essences*, 2013). Their mapping onto the admissibility schema is the authors' interpretive proposal.

4.1. Cognition as analogy-driven reduction

Hofstadter's central claim is that cognition is not primarily rule-following or symbolic deduction but *analogy-making*: concepts are "bundles of analogies" and thought proceeds through fluid transitions between them.

The process Hofstadter calls the *central cognitive loop* is structurally isomorphic to explicit substitution:

1. a stored conceptual structure (chunk) is activated;

2. partial unpacking into working memory occurs;
3. analogical perception propagates through the unpacked structure;
4. new conceptual activations emerge and are re-chunked.

Compare this to β-reduction:

1. a λ-abstraction is activated by application;
2. the argument is substituted through the scope;
3. local redexes are exposed and further reduced;
4. a new (possibly partially reduced) term emerges.

4.2. Chunking as nested bubble formation

Hofstadter's *chunking* operation recursively compresses conceptual structures into larger units. Small perceptual features become concepts; concepts become higher-order schema; schema become background assumptions that operate below the threshold of attention.

In Spherepop terms: each chunk is a bubble whose interior has been stably collapsed and whose membrane presents only a high-level interface to further interactions. Nested chunks are nested bubbles. Recursive chunking is the construction of deeply nested bubble hierarchies in which interior structure has been successively collapsed into stable sub-topologies.

Definition 4.1 (Cognitive admissibility)**.** A perceptual input Δ is *cognitively admissible* relative to a conceptual chunk $\mathbf{b}$ if Δ activates one or more attachment sites in $\partial(\mathbf{b})$ through analogical resonance, i.e., if a structural similarity mapping exists between Δ and components of $I(\mathbf{b})$.

4.3. Analogy, slippage, and fuzzy collapse

Hofstadter emphasizes that concepts are *blurry attractors*, not rigid symbolic containers. Analogical mapping is never exact: "slippage" occurs when a mapping activates a related but non-identical concept.

Standard λ-calculus is discrete and exact: substitution either succeeds or fails. Spherepop, by contrast, can in principle accommodate *partial* and *graded* admissibility: a bubble whose interior is only partially compatible with an incoming perturbation can undergo a *partial collapse* that leaves residual open sites. This is precisely the formal structure required for analogy-making, semantic slippage, contextual reinterpretation, and metaphorical extension.

Table 2: Three-way correspondence: λ-calculus, Hofstadter cognition, Spherepop.

λ-calculus	**Hofstadter cognition**	**Spherepop**
Variable	Open conceptual slot	Attachment site
Abstraction	Encapsulated procedure	Scoped bubble
Application	Conceptual activation	Bubble injection
β-reduction	Analogical rebinding	Membrane collapse
α-conversion	Context-sensitive reinterpretation	Historical relabeling
Normal form	Stable conceptual equilibrium	Stable bubble topology
Divergence	Recursive reminding loops	Infinite bubbling
Confluence	Semantic coherence	Topological reconciliation

4.4. Communication and semantic transport

Hofstadter describes language as *compressed symbolic transport*: a speaker encodes a complex meaning into a compact symbolic string; the listener reconstructs meaning by "adding water" — activating relevant conceptual structures and allowing them to partially re-expand.

In Spherepop: communication succeeds when the receiver's bubble topology is sufficiently compatible with the sender's to admit a successful membrane collapse. Translation across languages or conceptual frameworks becomes topology-preserving reduction across incompatible semantic manifolds, and failure to communicate is a gluing failure in the sense of Section 11.

4.5. Implication for Curry–Howard

Under the Spherepop–Hofstadter synthesis, the Curry–Howard correspondence acquires a cognitive reading:

- *types* are admissibility membranes (conceptual constraints);
- *proofs* are stabilized analogy trajectories;
- *programs* are constrained conceptual collapses;
- *normalization* is cognitive coherence formation.

Cognition is a generalized proof-reduction process operating over analogical topologies rather than purely syntactic strings.

5. Analogy as Reduction

Epistemic Status: Speculative Theoretical Proposal

The argument in this section is a theoretical proposal synthesizing Johansson's linguistic framework with the Spherepop–Hofstadter synthesis developed in the preceding sections. Johansson's specific claim that "analogy formations could be viewed as a process of reduction" is from his published work on construction and reduction in language acquisition. The generalization to a cross-domain admissibility-reduction schema is the authors' proposal.

5.1. Johansson's precedent

The present framework extends earlier reduction-based accounts of cognition and language by proposing that analogy itself is a reduction operation. This claim is not intended merely as metaphor.

Johansson's *Construction as Reduction* argues that language acquisition is not fundamentally constructive in the traditional sense but rather a cumulative process of selective reduction from a surrounding linguistic population. Crucially, Johansson explicitly states that "analogy formations could be viewed as a process of reduction," since analogy progressively eliminates irregularity and reduces variation across linguistic forms.

The present work generalizes this insight beyond linguistic evolution into a broader operational framework spanning computation, cognition, semantics, and physical systems.

5.2. Analogy as primary reduction mechanism

In ordinary symbolic models, analogy is treated as a secondary or heuristic process operating over already-defined structures. Under the present framework, analogy becomes a *primary* reduction mechanism through which structures themselves stabilize.

The essential claim: cognition survives precisely because it continuously reduces the number of admissible trajectories available to interpretation and action. Analogical perception performs this reduction by collapsing distinct configurations into families of structurally compatible transformations.

The analogical operator may be written:

$$\mathcal{A} : \mathcal{B}_1 \rightsquigarrow \mathcal{B}_2$$

where $\mathcal{A}$ reduces the distance between admissible structural configurations $\mathcal{B}_1$ and $\mathcal{B}_2$ by identifying a constraint-preserving mapping between their interiors.

This reformulates analogy with formal precision: *two structures are analogically related if a constrained transformation preserves enough invariant structure for coherent rebinding.* This is the admissibility-transport reading of analogy: vastly more rigorous than simply “finding similarities.”

5.3. Cross-domain instances

The same reduction grammar appears across all domains treated in this monograph:

- In **linguistic evolution**: analogy reduces morphological irregularity across a linguistic population;
- In **semantic cognition**: analogy reduces interpretive uncertainty by collapsing competing conceptual trajectories;
- In **RSVP field dynamics**: entropy gradients reduce admissible future states toward locally stable attractors;
- In **sheaf semantics**: successful gluing reduces local semantic incompatibilities into coherent global sections;
- In **Marine processing**: salience extraction reduces high-dimensional waveform variability into stable jitter structures;
- In **quantum interference**: coherent phase interaction reduces propagation histories into measurable amplitudes.

What recurs is not literal ontological identity but a shared operational grammar: *stable systems preserve admissibility through local reduction.*

5.4. Construction through recursive reduction

This yields a reversal of the traditional opposition between creativity and reduction. Construction may emerge through recursive reduction processes operating across populations, histories, and semantic topologies. Johansson’s formulation already hints at this inversion when language formation is described as a “stasis between constructive and eliminative reduction.”

The present framework extends that principle: reduction is not merely elimination. Properly understood, reduction is the mechanism through which coherence emerges from excess possibility.

This directly extends the monograph’s central thesis:

> **Analogy is admissibility-preserving transport. β-reduction is its symbolic special case. Membrane collapse is its geometric realization. Thermodynamic relaxation is its physical realization.**

Part II

Biological Instantiation: Multimodal Lineage Tracing

6. Cellular Differentiation as Admissibility-Constrained Collapse

Epistemic Status: Established (peer-reviewed results)

The biological material in this section is drawn from Wang, He, and Hu (2026), "Computational approaches for multimodal lineage tracing," *Nature Reviews Genetics*, 10.1038/s41576-026-00969-9. Specific characterizations of lineage-tracing methodology, ancestral state reconstruction, and computational frameworks follow that review. The interpretation of this material through the admissibility framework is the authors' proposal and does not represent the claims of the original authors.

6.1. The underdetermination of cellular identity by present state

In classical single-cell biology, cellular identity is treated as a point in a high-dimensional molecular state space, typically defined by a transcriptomic profile specifying the expression level of each gene in the cell's genome. Two cells with similar transcriptomic profiles are treated as similar cells; two cells with divergent profiles are treated as distinct. This framework is ontologically static: it describes what a cell is at the moment of measurement, not how it came to be what it is or what it can become.

The review by Wang, He, and Hu identifies a fundamental limitation of this static framework. Transcriptomic similarity does not reliably imply common developmental ancestry, because phenotypically similar cells may emerge through entirely different developmental histories. Conversely, cells that appear transcriptomically distinct may share recent common progenitors and therefore occupy the same admissibility basin with respect to future developmental potential. The present configuration of a cell underdetermines its generative history, its inherited constraint structure, and the space of its admissible futures.

This is a precise biological instance of the underconstraint problem developed in Section 18. A transcriptomic measurement generates a local section of the developmental manifold at a single moment, but the space of developmental trajectories compatible with that measurement is large. Multiple incompatible developmental histories may produce indistinguishable terminal states. Conversely, the same ancestral state may produce divergent terminal states depending on the history of environmental perturbations, signaling inputs, and stochastic gene-expression fluctuations encountered along the way. In the language of the admissibility framework, present configuration alone does not determine admissibility history, and admissibility history is what constrains admissible futures.

6.2. Lineage tracing as historical admissibility reconstruction

Multimodal lineage tracing addresses this limitation by coupling transcriptomic measurements with orthogonal records of developmental ancestry. These lineage records take several forms, including somatic mutations accumulated during cell division, CRISPR-induced barcodes introduced at defined developmental stages, DNA methylation patterns inherited across cell generations, and spatial position within tissue architecture. Each modality encodes a different aspect of the cell's developmental history, and their integration progressively constrains the space of developmental trajectories compatible with the observed present state.

From the perspective of the admissibility framework, lineage tracing is an attempt to reconstruct the sequence of admissibility regions $\mathcal{B}_0 \to \mathcal{B}_1 \to \cdots \to \mathcal{B}_n$ that generated the observed terminal configuration $\mathcal{B}_n$. Each transition in this sequence is a membrane collapse in the sense of Definition 11.1: a bounded developmental state accepts a compatible perturbation (a signaling input, a morphogen gradient, a mechanical force, a stochastic transcriptional fluctuation), internally reorganizes its constraint structure (chromatin remodeling, transcription factor redistribution, metabolic reprogramming), and collapses into a new admissible successor state while preserving certain invariants (cell viability, genomic integrity, tissue compatibility). The invariant that lineage tracing tracks is ancestral identity: the inherited sequence of developmental collapses that produced the current state.

The review characterizes this problem precisely when it describes multimodal lineage tracing as attempting to reconstruct "trajectory inference" and "ancestral state estimation" across the developmental manifold. The key challenge is that single-cell sequencing is destructive: the cell is lysed to extract its molecular contents, so the measured cell no longer exists after measurement. Ancestral states are therefore fundamentally unobserved, and computational systems must reconstruct them from the terminal-leaf observations alone, inferring which historical sequences of admissible collapses could have generated the observed distribution of present states.

6.3. Disparate multimodality and constraint reconciliation

A central technical challenge identified in the review is what the authors call "disparate multimodality": lineage relationships are discrete and tree-structured, while molecular profiles are continuous and high-dimensional. The lineage tree assigns each cell a discrete ancestry path through a bifurcating graph, while the transcriptomic state space assigns each cell a position in a continuous high-dimensional manifold. These two representations encode fundamentally different aspects of the cell's identity and are not immediately reconcilable by standard mathematical tools.

This disparate multimodality is, in the language of the admissibility framework, the problem of reconciling historical constraint topology with instantaneous manifold geometry. The lineage tree functions as a constrained admissibility skeleton: it records which developmental collapses have occurred in what order, and each branch point represents an irreversible commitment to a developmental trajectory that closes off other potential futures. The transcriptomic manifold is the continuous semantic field in which the consequences of those discrete commitments are expressed.

Formally, let T denote the lineage tree, with nodes representing cell generations and edges representing division events, and let $\mathcal{M}$ denote the transcriptomic manifold with points representing cell states. The reconstruction problem requires a map $\phi : T \to \mathcal{M}$ that is consistent with both the discrete ancestral relationships encoded in T and the continuous trajectory structure observed in $\mathcal{M}$. Such a map exists globally if and only if the local section assignments are pairwise compatible and satisfy the cocycle condition: precisely the sheaf-gluing condition of Appendix C. When the gluing condition fails, the lineage reconstruction is obstructed, and the system cannot produce a globally coherent developmental history consistent with all measurements simultaneously.

6.4. Developmental hallucination as gluing obstruction

The parallel with semantic hallucination identified in Section 11 and Appendix C is not merely formal. In both cases, local compatibility does not guarantee global coherence.

In the language model case, a system may produce locally fluent and pairwise-consistent text while generating a globally incoherent factual structure. The failure is detected only by checking triple overlaps rather than pairwise ones: the first Čech cohomology class $[\omega] \in \check{H}^1(X, \mathcal{F})$ is non-zero.

In the developmental biology case, the analogous failure occurs when locally plausible state assignments to individual cells cannot be assembled into a globally admissible lineage reconstruction. Two cells may each be individually consistent with their assigned transcriptomic profiles, and they may even be pairwise-consistent in the sense that their profiles are compatible with a common progenitor under a simple differentiation model, but the full reconstruction across all cells in the dataset fails because the implied ancestral trajectories violate the constraints imposed by the lineage tree structure.

This is what may be called *developmental hallucination*: a computational system assigns locally plausible cell states and cell transitions, but the resulting developmental history is globally inadmissible. It implies progenitor states that are biologically implausible, transition sequences that violate chromatin accessibility constraints, or ancestry assignments that contradict the discrete lineage topology. The hallucination is invisible at the level of individual cell-type assignments and becomes apparent only in the global reconstruction.

The review identifies optimal transport methods such as LineageOT and moslin as computational strategies specifically designed to address this problem. These methods seek the most economical mapping between developmental states under lineage constraints, minimizing admissibility violation while preserving inherited structural relationships. In the admissibility framework, these methods are approximating the globally consistent section of the developmental sheaf: the unique assignment of cells to developmental trajectories that satisfies both the local transcriptomic compatibility conditions and the global ancestral coherence requirements.

6.5. Progressive potency restriction as admissibility contraction

The review's concept of "progressive potency restriction" is one of the most structurally informative formulations in the paper from the perspective of the admissibility framework. As cells differentiate, they progressively lose the capacity to adopt alternative cell fates. A totipotent zygote can in principle give rise to any cell type. A pluripotent stem cell can give rise to most cell types. A committed progenitor can give rise to a restricted set of related cell types. A terminally differentiated cell typically cannot give rise to any other cell type at all.

This progression is not merely a statistical observation about which cell types tend to appear at which developmental stages. It is a constraint on the admissible future manifold of each cell. As differentiation proceeds, the set of admissible successor states contracts. Each developmental collapse irreversibly eliminates trajectories from the cell's future possibility space. The contraction is not merely a reduction in the number of available gene expression states but a reduction in the dimensionality of the admissibility region itself: the constraints that govern which transitions are biologically realizable become progressively tighter as inherited chromatin structure, DNA methylation patterns, and committed transcription factor networks accumulate.

In the formal language of Appendix D, this is entropy reduction in the specific sense of the RSVP framework: not thermodynamic entropy alone, but the logarithmic volume of admissible future trajectories. The entropy of an undifferentiated cell with respect to its developmental future is high; the entropy of a terminally differentiated cell is zero or near zero. Differentiation is the progressive reduction of this developmental entropy through irreversible admissibility collapse.

This interpretation is consistent with the review's observation that computational approaches to developmental dynamics often model cell differentiation through Waddington-landscape-style potential energy surfaces, in which cells occupy attractors separated by potential barriers. In the admissibility framework, these attractors are admissible equilibria of the developmental admissibility functional, the barriers are inadmissibility regions that require non-admissible perturbations to traverse, and the landscape itself changes over developmental time as the constraint structure inherited from prior collapses modifies the admissibility topology.

6.6. Deep learning as trajectory compression

The review's treatment of deep-learning frameworks for multimodal lineage analysis is philosophically significant because it illustrates how modern computational systems operationalize the admissibility framework implicitly. Methods such as LineageVAE, PORCELAN, and the deep lineage frameworks reviewed in the paper construct latent embeddings of cell state that are simultaneously constrained by transcriptomic similarity and lineage topology.

In the admissibility framework, these methods are learning compressed representations of the developmental admissibility manifold: lower-dimensional summaries that preserve the historical constraint structure required for coherent trajectory inference. The variational autoencoder objective in LineageVAE, for instance, balances a reconstruction term (fidelity to observed transcriptomic state) against a regularity term (prior structure on the latent space), while the lineage coupling introduces an additional constraint that latent representations of ancestrally related cells must lie in geometrically consistent regions of the latent manifold.

This is formally analogous to the CLIO projection framework introduced in Section 16. In CLIO, the observer navigates a compressed projection $\pi : X \to M$ from a high-dimensional ontic trajectory space X to an operational measurement manifold M, preserving only the dynamically relevant constraint relations. The deep lineage methods are learning precisely this kind of projection: a compression that preserves the admissibility-bearing constraint structure of the developmental history while discarding the high-dimensional details that are not required for trajectory inference.

The review notes that these methods remain limited by the heterogeneity and incompleteness of multimodal lineage data, and by the difficulty of learning latent representations that correctly balance the disparate topological structures of lineage trees and continuous transcriptomic manifolds. In the admissibility framework, this limitation corresponds to the difficulty of learning an admissible projection when the admissibility structure of the full developmental manifold is accessible only through sparse, destructive, and temporally incomplete measurements. The deep-learning system is attempting to infer $\mathcal{R}$ from observations of $\mathcal{B}_{i+1}$ alone, without access to the intermediate collapse operators or the full history $H(\mathcal{B}_{i+1})$.

6.7. Admissibility as biological necessity, not metaphorical decoration

The most important philosophical consequence of this biological grounding is that admissibility is here not a speculative theoretical overlay but a structural necessity of the system. Cellular differentiation cannot explore arbitrary molecular configurations. The space of biologically realizable developmental trajectories is a small fraction of the space of all conceivable transitions between molecular states, constrained by heredity, chromatin architecture, spatial geometry, signaling network topology, energy budget, and the irreversibility of committed developmental collapses.

Without admissibility constraints, multicellular coherence would dissolve. If cells could transition arbitrarily between molecular states, the coordination between cells required for tissue organization, organ formation, and systemic physiology would become combinatorially impossible. The organism survives precisely because developmental collapse is constrained: each cell's trajectory is bounded by inherited admissibility structure, and that structure propagates across cell generations through mechanisms of epigenetic inheritance, gap junction communication, morphogen gradients, and mechanical coupling.

The admissibility framework therefore finds in multimodal lineage tracing not an analogy but an instantiation. The admissibility regions are the developmentally accessible chromatin configurations. The perturbations are the signaling inputs and environmental fluctuations to which the cell is exposed. The compatibility conditions are the biochemical requirements for signal transduction. The collapse operators are the gene regulatory network dynamics that integrate the perturbation and reorganize the cell's transcriptional state. The invariant is cellular viability and genomic integrity. The normal form is the terminally differentiated cell type. The failure mode is developmental mis-specification, in which a cell collapses into an inadmissible state, typically resulting in cell death, oncogenic transformation, or developmental arrest.

The review by Wang, He, and Hu (2026) thus provides an empirically grounded, computationally active, and formally tractable domain in which the admissibility principle is not a philosophical hypothesis but a computational necessity: the reconstruction of developmental history requires admissibility constraints, and those constraints are mathematically specifiable, biologically motivated, and progressively recoverable through the integration of multiple heterogeneous data modalities.

Part III

Established Physics

7. Quantum Weak-Value Dwell Times

Epistemic Status: Established (peer-reviewed results)

Results in this section are drawn from: APL Quantum **2**, 036108 (2025) on atomic excitation times for transmitted and scattered photons. Statements attributed to that paper reflect its conclusions; the admissibility mapping is the authors' interpretation.

7.1. The photon dwell-time problem

When a single photon propagates through a cloud of two-level atoms, it may briefly exist as a collective atomic excitation (a Dicke-like polariton state) before being re-emitted. The question of how long the photon "spends" as an atomic excitation has a physically non-trivial answer.

For *transmitted* photons, the atomic excitation time τ_T equals the net group delay:

$$\tau_T = \tau_{\text{group}}.$$

In resonant regimes where dispersion produces negative group delay (the "superluminal" regime), τ_T is negative. This is not a violation of causality but an instance of an *anomalous weak value*: a postselection effect arising from quantum interference between two pathways.

For *scattered* photons, the excitation time involves an additional contribution:

$$\tau_S = \langle \tau_{\text{group}} \rangle_{\text{ens}} + \tau_{\text{Wigner}},$$

where τ_{Wigner} is the Wigner time delay associated with elastic scattering.

7.2. Anomalous weak values as interference-mediated admissibility

The negative-time result arises from interference between:

1. a pathway in which the photon passes without excitation, and
2. a pathway in which the photon converts to an atomic excitation and coherently transfers back into the photonic mode.

Destructive interference between these pathways produces a net negative pointer shift — an anomalous weak value.

Note 7.1 (Caveat on "negative time"). The phrase "the photon spends a negative amount of time as an excitation" is shorthand for a postselected weak-value measurement. It does not imply backward causation. The weak value is a property of the measurement protocol (pre- and post-selection), not a directly observable duration in the classical sense.

7.3. Admissibility mapping

The quantum dwell-time experiment instantiates the admissibility schema as follows:

- **Admissibility region** $\mathcal{B}_i$: the atomic medium with its resonance structure and spectral response function;
- **Perturbation** Δ_i: the incoming photon wavepacket;
- **Constraint / compatibility test**: the photon's spectral components must overlap the atomic resonance for excitation to be possible; the interference condition determines admissibility of the excitation pathway;
- **Collapse operator** $\mathcal{R}$: destructive/constructive interference between pathways;
- **Stable output** $\mathcal{B}_{i+1}$: transmitted photon with modified group delay, or scattered photon with Wigner delay.
- **Invariant preserved**: unitarity / probability conservation.

The confluence principle here is the Born rule: different measurement protocols on the same quantum state must produce statistically consistent results.

8. DC-OPA and Ultrafast Nonlinear Optics

Epistemic Status: Established (peer-reviewed results)

Results in this section are drawn from the advanced DC-OPA experiment reported in *Nature Photonics* (2024–25): a dual-chirped optical parametric amplification system using BiBO and MgO:LN crystals. Specific figures (pulse duration, peak power, CEP stability) are taken from that source. Claims about water-window harmonic generation efficiency and attosecond pulse duration are the authors' reading of the literature and should be independently verified.

8.1. Dual-chirped optical parametric amplification

High-order harmonic generation (HHG) for attosecond science requires few-cycle mid-infrared (MIR) drivers. Standard OPCPA (optical parametric chirped-pulse amplification) stretches a seed pulse, amplifies it parametrically, and compresses it back to short duration.

Dual-chirped OPA (DC-OPA) goes further: both the pump *and* the seed are chirped. This relaxes phase-matching bandwidth constraints and allows amplification over more than one octave using heterogeneous nonlinear crystals (BiBO for the first stage; MgO-doped lithium niobate for the second).

The reported performance of the advanced DC-OPA system:

- Output energy: 53 mJ;
- Pulse duration: 8.58 fs at 2.44 μm central wavelength (≈ 1.05 optical cycles);
- Peak power: $\sim$6 TW;
- CEP stability: 228 mrad.

8.2. Phase-matched energy redistribution as constrained collapse

The physics of parametric amplification is governed by phase-matching: energy transfer from pump to signal and idler is efficient only when the wavevectors satisfy

$$\Delta k = k_p - k_s - k_i = 0.$$

The nonlinear crystal is the constraint membrane; the incoming pump and seed are the perturbation; phase-matched energy redistribution is the collapse operation; the output compressed pulse is the stable attractor.

- **Admissibility region $\mathcal{B}_i$**: the nonlinear crystal with its phase-matching bandwidth and group-velocity dispersion profile;
- **Perturbation Δ_i**: the chirped pump and seed pulses;
- **Compatibility condition**: $\Delta k \approx 0$ over the amplification bandwidth;
- **Collapse operator $\mathcal{R}$**: nonlinear three-wave mixing under phase-matching constraint;
- **Stable output $\mathcal{B}_{i+1}$**: compressed MIR pulse with coherent phase structure;
- **Invariant**: photon number conservation (Manley–Rowe relations); CEP coherence.

8.3. HHG and attosecond science

The compressed MIR pulses drive HHG in gas targets. An electron is ionized by the intense field, accelerated, and recombines with the ion, emitting a burst of XUV or soft X-ray radiation. The three-step model (ionization, propagation, recollision) is itself an admissibility sequence: recollision is gated by the phase structure of the driving field, and only electron trajectories that satisfy the recollision condition contribute to harmonic emission.

The extension to isolated attosecond pulses (IAPs) requires spectral selection from a supercontinuum region where contributions from multiple half-cycles do not interfere destructively. This is a second-level admissibility gate on top of the first.

Part IV

Speculative and Prototypical Layers

9. RSVP: Thermodynamic Field Semantics

Epistemic Status: Speculative Theoretical Proposal

The Relativistic Scalar-Vector Plenum (RSVP) framework is a theoretical proposal under development by the author. Parameters such as the critical coupling $\lambda_c \approx 0.42$ and the entropy production threshold $\dot{\Sigma}_{\mathrm{crit}} = 2.1$ nats/generation are *model-derived conjectural thresholds*, not empirically measured physical constants. They should not be read as established results of the same evidential status as the quantum-optics or ultrafast-laser material in Part II.

9.1. Framework overview

RSVP models the evolution of physical and cognitive systems through a coupled field theory over three interacting fields:

- Φ: a scalar entropy-potential field;
- $\mathbf{v}$: a vector flow field encoding directed processes;
- S: an entropy density field.

The central postulate is that admissible system trajectories are those along which entropy production remains bounded: $\dot{S} \geq 0$ locally (second law), but $\dot{S} \leq \dot{\Sigma}_{\mathrm{crit}}$ globally for stable operation (stability condition).

9.2. RSVP as thermodynamic admissibility collapse

In the RSVP framework, a local region of the field undergoes admissibility-constrained collapse as follows:

- **Admissibility region $\mathcal{B}_i$**: a connected region of $(\Phi, \mathbf{v}, S)$ field space with locally defined entropy gradient ∇S;

- **Perturbation Δ_i**: an incoming entropy flux or directed flow perturbation;
- **Compatibility condition**: the perturbation must be absorbable without driving $\dot{S}$ beyond $\dot{\Sigma}_{\text{crit}}$;
- **Collapse operator $\mathcal{R}$**: entropy-gradient-constrained field relaxation;
- **Stable output $\mathcal{B}_{i+1}$**: a new field configuration with lower free energy and admissible entropy production;
- **Invariant**: total entropy non-decrease; field continuity.

9.3. The thermodynamic reading of β-normal form

In λ-calculus, a β-normal form is a term with no remaining redexes: computation has reached a stable configuration. In RSVP, the corresponding notion is a field state in which no admissible entropy-decreasing trajectory remains locally available: a thermodynamic equilibrium or metastable attractor.

The analogy is not that computation and thermodynamics are the same physics, but that both are instances of the generalized reduction schema in which "no further collapse is possible" constitutes the normal form.

9.4. RSVP parameters as scaling hypotheses

The numerical thresholds $\lambda_c \approx 0.42$ and $\dot{\Sigma}_{\text{crit}} = 2.1$ nats/generation are *provisional scaling hypotheses* defining admissible behavioral regimes within RSVP simulation work. They are not empirical constants, not derived from first principles within an established physical theory, and not published in peer-reviewed venues.

The appropriate analogy is Reynolds-number thresholds in early turbulence research: numbers that organized experimental intuition and defined qualitative regime boundaries before quantitative derivation was available. The RSVP constants should be treated as heuristic bifurcation indicators subject to revision as the framework is developed.

9.5. Thermodynamic constraints versus impossibility proofs

The RSVP framework motivates a thermodynamic argument against unbounded recursive self-improvement ("FOOM"):

1. Any cognitive subsystem increasing organizational complexity must export entropy to its environment (second law applied to open systems).
2. The entropy export rate is bounded by thermal conductivity and effective surface area.
3. Therefore, the rate of self-improvement is bounded by a finite thermodynamic flux.

This argument is a *scaling constraint*, not a hard impossibility theorem. Several mechanisms could alter the bounds:

- *Reversible computing*: Landauer's principle permits logically reversible operations to approach zero heat dissipation per step; a sufficiently reversible architecture could reduce entropy export per operation arbitrarily.
- *Quantum adiabatic evolution*: adiabatic processes produce no entropy in the ideal limit.
- *Distributed extension*: a system growing spatially can increase effective surface area, relaxing the area bound.

The correct claim is: *classical, locally bounded self-improvement faces finite thermodynamic constraints under the RSVP parameters as currently defined.* Whether those constraints bind depends on architecture and regime. This is a philosophical motivation and a modeling choice, not a theorem of the status of the second law itself.

10. Marine, Phoenix 8a, and MEM|8

Epistemic Status: Engineering Prototype / Proposed Architecture

Phoenix 8a, Marine pattern memory, and MEM|8 are engineering architectures at the prototype or proposal stage. Benchmarks quoted (RMS match accuracy $> 99.5\%$, voice cloning from minimal data) are internal test results that have not been independently replicated or published in a peer-reviewed venue. Claims about "infinite frequency resolution" and "zero latency" are significant overstatements of information-theoretic limits and must be qualified (see Section 10.3). The material here is presented as motivated engineering hypothesis, not established science.

10.1. The core design philosophy

Traditional digital audio processing converts continuous waveforms into sampled pulse-code modulation (PCM) and analyzes them via the Short-Time Fourier Transform (STFT) or Mel spectrogram. The fundamental trade-off is the time-frequency uncertainty relation:

$$\Delta t \cdot \Delta f \geq \frac{1}{4\pi}.$$

Large analysis windows provide high frequency resolution but poor time resolution, destroying attack transients and micro-temporal features that carry perceptual information about articulation, emotional valence, and speaker identity.

The Marine engine proposes an alternative: *time-domain interval counting.* Rather than sampling waveform amplitude at a fixed rate and Fourier-transforming, Marine tracks the intervals between zero-crossings and local extrema, extracting pitch, ADSR envelopes, vibrato trajectories, and jitter signatures directly from the temporal structure of the waveform.

10.2. Admissibility mapping

- **Admissibility region $\mathcal{B}_i$**: the temporal-interval structure of a waveform segment, interpreted as a constraint membrane over the space of admissible harmonic families;
- **Perturbation Δ_i**: an incoming waveform event (zero-crossing, peak, onset);
- **Compatibility condition**: the event's temporal position must be consistent with the current harmonic-family hypothesis;
- **Collapse operator $\mathcal{R}$**: jitter/salience extraction, ADSR envelope estimation, harmonic-lock detection;
- **Stable output $\mathcal{B}_{i+1}$**: a compressed representation as a sparse set of harmonic families with associated envelope and jitter parameters;
- **Invariant**: temporal ordering; pitch contour continuity.

10.3. Caveat: time-frequency uncertainty

The phrase "infinite frequency resolution with zero latency" cannot be taken literally: it violates the time-frequency uncertainty principle, which is not an artifact of FFT implementation but a consequence of the Fourier duality between time and frequency representations. What the Marine approach *can* legitimately claim is:

- resolution of *pitch* from temporal periodicity without a fixed window function;
- lower latency than STFT for onset detection;
- extraction of temporal features (jitter, vibrato rate) that STFT discards by design.

These are meaningful advantages that should be claimed on their own terms rather than through a physically incoherent slogan.

10.4. MEM|8 as persistent attractor topology

MEM|8 proposes to store compressed audio "atoms" (harmonic families and residual burst structures) as activation seeds for a semantic memory system. The design goal is that memory retrieval proceeds by *analogical resonance* rather than hash-lookup:

a query activates the stored atoms whose structure is most compatible with the query's temporal-harmonic profile.

In Hofstadter's terms, this is a formal implementation of analogy-driven memory reconstruction. In Spherepop terms, stored atoms are bubble-topology snapshots; retrieval is admissibility-gated membrane activation. Whether current implementations of MEM|8 realize this design goal at the claimed fidelity levels is a question that requires independent benchmark evaluation.

11. The Generalized Reduction Operator

11.1. Formal statement

We now bring together the preceding analyses into a single schema.

Definition 11.1 (Generalized reduction operator)**.** Let $\mathcal{B}$ be a class of *bounded admissibility regions* and $\mathcal{D}$ a class of *perturbations.* A *generalized reduction operator* is a partial function

$$\mathcal{R} \,:\, \mathcal{B} \times \mathcal{D} \rightharpoonup \mathcal{B}$$

satisfying:

(i) **Admissibility gate**: $\mathcal{R}(\mathcal{B}_i, \Delta_i)$ is defined only when Δ_i is compatible with the constraint structure $\partial(\mathcal{B}_i)$;

(ii) **Invariant preservation**: there exists an invariant function $\iota : \mathcal{B} \to I$ such that $\iota(\mathcal{R}(\mathcal{B}_i, \Delta_i)) = \iota(\mathcal{B}_i)$ whenever $\mathcal{R}$ is defined;

(iii) **Irreversibility**: the history H is strictly extended by each application of $\mathcal{R}$;

(iv) **Locality**: $\mathcal{R}(\mathcal{B}_i, \Delta_i)$ depends only on the local constraint structure $\partial(\mathcal{B}_i)$ and the interior $I(\mathcal{B}_i)$, not on the global state of the system.

11.2. Necessary and sufficient conditions for instantiating $\mathcal{R}$: what does *not* qualify

The reviewer of any draft of this work will correctly ask: *what would not count as an instance of* $\mathcal{R}$? Without a nontrivial exclusion criterion, the schema risks vacuity — any constrained process with inputs and outputs could be force-fit into the table. This subsection addresses that objection directly.

Three necessary conditions beyond Definition 11.1. A process qualifies as an admissibility-collapse system only if all four conditions of Definition 11.1 hold *and* three additional structural requirements are met:

1. **Recursive state-space restructuring.** Admissibility must actively change the future trajectory space, not merely filter outputs after the fact. A thermostat does not qualify: it enforces a temperature band but does not restructure the space of admissible future states. A lookup table does not qualify: it maps inputs to outputs without recursively reorganizing the constraint structure. The test is whether removing the admissibility gate would allow the system to reach qualitatively different macrostates.

2. **Transportable invariant.** Collapse must propagate a coherent invariant through the transformation, not merely destroy structure. A shattered glass is not performing admissibility-guided reduction: no coherent invariant (ι) survives the transformation in a form that constrains future evolution. The invariant must be *transportable* — usable by downstream processes as a constraint.

3. **Path-sensitive historical dependence.** The system's admissibility structure must depend on its reduction history, not merely its current state. This is condition (iii) of Definition 11.1 made operationally precise. A Markov process with no memory of prior states does not qualify unless its transition kernel is itself modified by prior reductions.

Canonical non-instances. The following processes explicitly do not instantiate $\mathcal{R}$:

- *Thermostats and bang-bang controllers*: constraint without state-space restructuring.
- *Hash lookups and lookup tables*: output filtering without invariant propagation.
- *Markovian random walks*: locally defined transitions without historical dependence.
- *Physical destruction* (shattering, combustion without product reuse): transformation without transportable invariant.
- *Arbitrary compression* (e.g., lossy JPEG): projection without admissibility gate or invariant preservation.

Why biological development qualifies. A fertilized egg undergoing development is a canonical positive instance. Each admissible developmental collapse (cell fate determination, tissue boundary formation, organogenesis) recursively restricts future developmental possibilities while preserving organismic continuity. The space of possible futures narrows through a sequence of history-dependent admissibility collapses, each propagating structural invariants (positional information, morphogen gradients, lineage identity) that constrain subsequent stages. This is precisely the intended meaning of $\mathcal{R}$.

11.3. Domain specializations

Table 3 summarizes how each domain instantiates Definition 11.1.

Table 3: Concept atlas: domains as specializations of $\mathcal{R}$.

Domain	**Admissibility region $\mathcal{B}_i$**	**Perturbation Δ_i**	**Collapse operator $\mathcal{R}$**	**Invariant / failure mode**
λ-calculus	Scoped abstraction $\lambda x.\, M$	Argument N	β-reduction: $M[x := N]$	β-normal form / divergence (Ω)
Spherepop	Scoped bubble with membrane ∂	Compatible bubble $\mathbf{b}'$	Membrane collapse $\oplus$	Stable topology / infinite bubbling
Hofstadter cognition	Conceptual chunk (bundle of analogies)	Perceptual input or memory cue	Analogical rebinding (central cognitive loop)	Stable concept / recursive reminding loop
Quantum dwell time	Atomic medium (resonance + dispersion)	Photon wavepacket	Interference-mediated pathway collapse	Group delay (possibly anomalous weak value) / decoherence
DC-OPA / HHG	Nonlinear crystal (phase-matching envelope)	Chirped pump + seed pulses	Three-wave mixing under $\Delta k = 0$	Compressed CEP-stable pulse / group-velocity mismatch
RSVP	$(\Phi, \mathbf{v}, S)$ field region	Entropy flux δS	Gradient-constrained field relaxation	Thermodynamic attractor / supercritical runaway
Marine / Phoenix 8a	Temporal-interval structure of waveform	Onset event (zero-crossing, peak)	Jitter + ADSR extraction, harmonic-lock	Sparse harmonic family representation / lock failure

continued on next page

Table 3 continued.

Domain	Admissibility region	Perturbation	Collapse operator	Invariant / failure mode
Sheaf semantics	Open cover of semantic patches	Overlap data between patches	Gluing map	Global section (coherent meaning) / gluing obstruction

11.4. Why confluence appears across domains: family resemblance, not a universal theorem

Note 11.2 (Terminological precision). The phrase “master convergence principle” used in earlier drafts of this work is *heuristic terminology*, not a proved universal theorem. The Church–Rosser theorem is a precisely stated and proved result about λ-calculus. The Born-rule consistency condition is a physical postulate about quantum measurement. The existence of global sheaf sections is a categorical result in algebraic topology. These are mathematically distinct statements that cannot be collapsed into a single theorem without a category-theoretic framework (e.g., showing that all are instances of limits preserving coherence in a 2-category) that is not provided here.

What can be legitimately claimed is that these domains exhibit a *family resemblance* generated by a common structural motif: *local compatibility reconciliation producing globally coherent output.*

Consider the analogical reading of each convergence result:

- In λ-calculus: incompatible substitution orders are reconciled through normalization to a common normal form.
- In sheaf theory: overlapping local descriptions are reconciled through gluing maps into a global section.
- In turbulence: incompatible local energy-transfer rates are reconciled through admissible dissipation into a single physically correct trajectory.
- In game theory: incompatible strategy profiles are reconciled through iterated dominance elimination into a common solution concept.
- In cognition: competing analogical interpretations are reconciled through contextual rebinding into stable conceptual equilibria.

The common element is not a shared theorem but a shared *question*: when a system faces locally incompatible pressures, does it possess a coherent collapse pathway? For λ-calculus the answer is the Church–Rosser theorem. For the other domains, confluence is a hypothesis, an engineering desideratum, or a phenomenon to be explained.

The admissibility principle is the claim that this question has the same *form* across domains, not that it has the same *answer*. Recognizing the shared form is the contribution; proving or disproving confluence in each domain is the research program.

11.5. Sheaf obstruction and hallucination

An important failure mode of admissibility-constrained collapse is *gluing obstruction*: local patches are each internally consistent but cannot be assembled into a globally coherent state.

In sheaf theory: a presheaf satisfies local compatibility but fails to admit a global section; the obstruction is measured by the first Čech cohomology group $\check{H}^1$.

In cognitive terms: individual analogical fragments activate without contradiction locally, but their combination produces a globally incoherent representation. This is a formal model of semantic hallucination in language models: the system generates locally plausible token sequences whose global factual structure violates world-knowledge constraints.

Proposition 11.3 (Hallucination as gluing failure)**.** *Let* $\{U_i\}$ *be an open cover of a semantic space, and let* $s_i \in \mathcal{F}(U_i)$ *be locally consistent semantic patches. A* hallucination *occurs when the cocycle condition* $s_i|_{U_i \cap U_j} = s_j|_{U_i \cap U_j}$ *fails for some* i, j*, so that no global section* $s \in \mathcal{F}(X)$ *restricting to all* s_i *exists.*

Part V

Convergent Literature and the Five-Layer Architecture

12. Game-Theoretic Admissibility: Recursive Elimination Under Common Knowledge

Epistemic Status: Established (peer-reviewed results)

The material in this section follows Bicchieri and Schulte (1996), "Common Reasoning About Admissibility," *Erkenntnis* (revised 1996). The mapping onto the Spherepop and sheaf frameworks is the authors' interpretive proposal.

12.1. Admissibility as a dynamical operator on possibility spaces

Bicchieri and Schulte introduce a dimension of admissibility absent from purely structural accounts: admissibility is not a static property of states but a *recursive elimination* process operating under common knowledge constraints.

Their central result concerns strategic games. A strategy s_i for player i is *(weakly) admissible* if it is not weakly dominated: no alternative strategy yields a better outcome against every possible opponent strategy. *Common reasoning about admissibility* is the iterated process by which players eliminate inadmissible strategies under the assumption that all players are doing likewise.

Definition 12.1 (Sequential proper admissibility)**.** A strategy s_i is *sequentially properly admissible* in a game tree T if s_i is admissible at every information set I_i consistent with s_i, i.e., at every node reachable under s_i.

The main theorem of Bicchieri and Schulte establishes that in finite extensive games with perfect recall, the strategies consistent with common reasoning about sequential proper admissibility in the extensive form coincide exactly with those consistent with common reasoning about admissibility in the strategic form. In other words: the solution given by recursive admissibility elimination is representation-independent.

12.2. Admissibility as trajectory pruning

The operational schema of Bicchieri–Schulte is:

possible trajectories $\longrightarrow$ eliminate inadmissible branches $\longrightarrow$ recurse under common knowledge

This maps directly onto our generalized reduction operator:

- **Admissibility region** $\mathcal{B}_i$: the information set of player i, together with the probability assessment over opponent strategies;
- **Perturbation** Δ_i: a new strategy or deviation possibility;
- **Compatibility condition**: the strategy must survive iterated weak-dominance elimination under strict coherence (every opponent strategy receives positive probability);
- **Collapse operator** $\mathcal{R}$: iterated elimination of weakly dominated strategies under common knowledge;
- **Stable output**: the set of rationalizable, admissibility- consistent strategy profiles;
- **Invariant**: representation-independence (strategic-form and extensive-form admissibility coincide).

The key philosophical point is that admissibility here is not static. It propagates globally through nested epistemic embeddings. This is structurally identical to:

- recursive sheaf gluing across overlapping open sets;
- Hofstadter's strange loops (cognition reasoning about its own reasoning);
- RSVP/CLIO recursive coherence filtering.

Table 4: Correspondence: game-theoretic admissibility and Spherepop.

Game theory (Bicchieri–Schulte)	**Spherepop**
Strategy	Bubble trajectory
Information set	Local membrane state
Weak dominance	Inadmissible collapse path
Iterated admissibility (IWD)	Recursive membrane pruning
Sequential admissibility	Path-consistent topology
Common knowledge of rationality	Shared admissibility manifold
Backward induction	Bottom-up bubble resolution
Forward induction	Top-down admissibility propagation

12.3. Backward and forward induction as admissibility projections

Bicchieri and Schulte show that backward and forward induction principles, often treated as mutually inconsistent, both emerge from one underlying principle: common reasoning about admissibility. Their unification parallels the monograph's broader claim that apparently distinct reasoning paradigms (symbolic, topological, thermodynamic, cognitive) are specializations of a common schema.

This also supplies the epistemic layer that earlier sections lacked. RSVP provides thermodynamic admissibility. Spherepop provides topological admissibility. λ-calculus provides syntactic admissibility. Bicchieri–Schulte provides *epistemic admissibility under recursive mutual constraint propagation.* The four layers form a complete stack.

13. Turbulence Admissibility: Physical Realization Requires Selection Beyond Dynamics

Epistemic Status: Established (peer-reviewed results)

Results in this section follow Glimm, Cheng, Sharp, and Kaman, "A crisis for the verification and validation of turbulence simulations," published preprint (Los Alamos LA-UR-19-20285). The three-algorithm comparison and maximum-dissipation validation evidence are as reported there. The admissibility mapping is the authors' interpretation.

13.1. The nonuniqueness problem in fluid dynamics

The Euler equations for incompressible fluid flow,

$$\partial_t \mathbf{v} + (\mathbf{v} \cdot \nabla)\mathbf{v} = -\nabla p, \quad \nabla \cdot \mathbf{v} = 0,$$

admit multiple physically distinct weak solutions even for identical initial conditions. This is not a numerical artifact: de Lellis and Székelyhidi have constructed infinitely many energy-conserving wild solutions. An additional physical principle is therefore required to select the physically realized trajectory among mathematically admissible ones. Glimm et al. call this the *admissibility principle for fluid dynamics.*

13.2. Three admissibility criteria and their consequences

Three physically motivated selection principles produce markedly different simulations of Rayleigh–Taylor (RT) turbulent mixing:

1. **Maximum dissipation rate** (FronTier algorithm): the local energy dissipation rate is maximized consistent with turbulent scaling laws.
2. **Limited dissipation** (ILES / Miranda): the Reynolds stress is treated as a Gibbs-type numerical artifact to be minimized.
3. **Near-zero dissipation** (MDNS): sub-grid-scale terms are suppressed based on a globally averaged Reynolds number.

The three algorithms produce RT mixing rates α_b that differ by factors of 2–3. Validation against experimental data from Smeeton–Youngs supports the maximum-dissipation principle.

The critical insight for the present framework:

$$\text{governing equations} + \text{admissibility principle} \longrightarrow \text{physically realized solution.}$$

Dynamics alone are insufficient.

13.3. Maximum entropy production as admissibility selector

The paper identifies maximum entropy production as the admissibility criterion consistent with experiment:

$$\dot{S} = \dot{S}_{\max}(x, t)$$

locally, under the weak quasi-stationarity hypothesis. This is a physical analogue of RSVP's thermodynamic admissibility: in both cases, entropy production is not merely a consequence of dynamics but a *constraint selecting* which among many dynamically consistent trajectories is realized.

Note 13.1 (Reynolds stress as admissibility-bearing structure). The Reynolds stress $\overline{uv} - \bar{u}\bar{v}$ is not numerical noise. It encodes genuine sub-grid-scale admissibility information that cannot be projected away without altering the macroscopic trajectory. Discarding it (ILES/MDNS) produces physically incorrect mixing rates. This is a turbulence analogue of the semantic compression argument: projections that discard constraint-carrying structure alter admissibility and produce incorrect global behavior.

13.4. Admissibility mapping for turbulence

- **Admissibility region $\mathcal{B}_i$**: the flow field at resolved scales with its local turbulent statistics;
- **Perturbation Δ_i**: sub-grid-scale Reynolds stress contributions;
- **Compatibility condition**: energy transfer must satisfy locally maximal dissipation under turbulent scaling laws;
- **Collapse operator $\mathcal{R}$**: dynamic sub-grid-scale model (FronTier-style) realizing maximum local dissipation;
- **Stable output**: physically correct RT mixing rate $\alpha_b \approx 0.06$;
- **Invariant**: energy conservation (Manley–Rowe at the fluid level); Kolmogorov scaling.
- **Failure mode**: inadmissible suppression of Reynolds stress → $\alpha_b \approx 0.02$–0.03, contradicting experiment.

14. Observer Admissibility in Relativistic Physics

Epistemic Status: Speculative Theoretical Proposal

The Modgil–Patil (2026) paper on observer admissibility and Planck invariance is not yet peer-reviewed in its current form and contains acknowledged mathematical gaps (particularly in the covariant formulation of γ_g and the phenomenological global-bound formula). The Evangelista (preprint) and Zenteno (Zenodo, 2026) papers are similarly pre-publication. The material in this section is presented as a convergent theoretical program, not established physics. The mapping onto the admissibility schema is the authors' synthesis.

14.1. The observerhood restriction strategy

A cluster of recent proposals converges on a common move: rather than modifying the dynamical laws of relativity to accommodate Planck-scale physics, restrict the class of physically realizable observers.

The key thesis, stated in Modgil–Patil, is:

> Planck-scale limits can coexist with exact local Lorentz invariance when they are interpreted as operational bounds on admissible observers, rather than as indicators of symmetry breaking or fundamental spacetime discreteness.

An observer is locally admissible when

$$g_{\mu\nu} V^\mu V^\nu > 0,$$

i.e., the observer remains timelike relative to the local metric. Global admissibility adds cosmological constraints:

$$\Delta t \leq T, \quad \Delta s \leq L, \quad m \leq M,$$

where T, L, M are finite characteristic scales of the universe. The physically realizable domain is the intersection:

$$\mathcal{A} = \mathcal{A}_{\text{local}} \cap \mathcal{A}_{\text{global}}.$$

14.2. Admissibility mapping for relativistic observerhood

- **Admissibility region** $\mathcal{B}_i$: the local spacetime neighborhood, together with global cosmological scale constraints;
- **Perturbation** Δ_i: a candidate measurement protocol (localization attempt, boost, long-duration process);
- **Compatibility condition**: $g_{\mu\nu} V^\mu V^\nu > 0$ locally; $\Delta t \leq T$, $\Delta s \geq \ell_P$ globally;
- **Collapse operator** $\mathcal{R}$: kinematic phase-boundary crossing (horizon encounter, Planck-scale localization attempt);
- **Stable output**: measurement result for an admissible observer; black-hole thermodynamics recovered without Lorentz violation;
- **Invariant**: exact local Lorentz symmetry; causal structure;
- **Failure mode**: inadmissible observer — stationary observer at horizon, trans-Planckian localization attempt, infinite boost.

14.3. The Kretschmann scalar as minimal admissibility detector

Evangelista (preprint) provides a rigorous geometric instantiation of the admissibility idea. He imposes three requirements on an admissibility obstruction functional $\mathcal{K}_{\mathcal{A}}[\Phi]$:

1. local and covariant;
2. scalar (observer-independent);
3. non-degenerate in Ricci-flat (vacuum) spacetimes.

The third condition is decisive: in Ricci-flat spacetimes, $R = 0$ and $R_{\mu\nu}R^{\mu\nu} = 0$, so lower-order scalars vanish. The *Kretschmann scalar*

$$K = R_{\mu\nu\rho\sigma}R^{\mu\nu\rho\sigma}$$

is the minimal quadratic invariant satisfying all three requirements: it detects purely tidal (Weyl) curvature even in vacuum.

Admissibility may then be expressed as the curvature bound

$$\hat{K} = K\ell_o^4 \leq 1,$$

which excludes curvature singularities from the physically realizable domain.

Remark 14.1 (Kretschmann as non-degenerate admissibility detector)*.* The Kretschmann scalar is a physical instance of a general principle that recurs throughout this monograph: *admissibility operators must remain sensitive to latent instability modes invisible to coarse observables.* In general relativity, lower-order Ricci scalars vanish in vacuum while dangerous Weyl curvature persists. Analogously, in language models, syntactic coherence fails to detect semantic hallucination; in ILES turbulence simulations, averaged Reynolds numbers fail to detect locally dangerous energy concentrations. The Kretschmann scalar formalizes what a non-degenerate admissibility detector must look like.

14.4. Zenteno: admissibility as dynamical stability criterion

Zenteno (2026) formalizes admissibility through continuous penalty functionals relative to constraint manifolds. A structure is admissible if it exhibits *bounded, stable, and resolvable* behavior under soft constraint dynamics. This generalizes classical existence-and-consistency criteria to include dynamical persistence:

Definition 14.2 (Zenteno admissibility)**.** A mathematical structure $\mathcal{S}$ is *admissible* relative to a constraint manifold $\mathcal{C}$ if the penalty functional

$$\mathcal{P}[\mathcal{S};\mathcal{C}] < \infty$$

and the trajectory $\mathcal{S}(t)$ under constraint dynamics remains within a bounded neighborhood of $\mathcal{C}$ for all t.

This framework naturally explains singularity exclusion: singular geometries violate the boundedness condition rather than the existence or consistency conditions. The Paton (2026) pre-explanatory admissibility layer makes the same move at the level of scientific modelling: a system state must satisfy structural membership conditions before any explanatory dynamics can meaningfully apply.

15. The Five-Layer Architecture

The preceding sections have accumulated sufficient material to articulate the monograph's full epistemic architecture explicitly. We formalize five layers, each building on the previous:

Layer I: Established empirical systems. Quantum optics (weak-value dwell times, APL Quantum 2025), ultrafast photonics (DC-OPA, Nature Photonics), gauge theory (standard), Rayleigh–Taylor turbulence validation (Glimm et al.). These are peer-reviewed results with experimental support. The admissibility schema maps onto them but is not required to explain them; they serve as *test cases* for the schema's descriptive coherence.

Layer II: Engineering architectures and computational hypotheses. Marine temporal-interval analysis, Phoenix 8a, MEM|8, wave-memory semantic retrieval. These are prototype-stage proposals requiring independent benchmark validation. Their connection to the admissibility schema is motivational, not evidentiary.

Layer III: Formal abstraction layer. β-reduction, sheaf gluing and obstruction cohomology, Church–Rosser confluence, game-theoretic iterated admissibility (Bicchieri–Schulte), HoTT transport, explicit substitution calculi. These are established mathematical frameworks. The generalized reduction operator $\mathcal{R}$ (Definition 11.1) lives at this layer.

Layer IV: Philosophical synthesis. Hofstadterian analogy and chunking, Spherepop membrane collapse as geometric operational semantics, RSVP thermodynamic cognition, observer admissibility as meta-constraint on coherent existence (Modgil–Patil, Zenteno, Paton). These are interpretive proposals organizing the lower layers. The admissibility principle as stated in this monograph lives primarily here.

Layer V: Research program. Define generalized reduction operators with explicit soundness proofs; develop admissibility metrics for each domain; investigate category-theoretic transport structures (natural transformations, fibered categories) capable of relating admissibility conditions across layers; formalize the generalized confluence schema; calibrate RSVP parameters from first principles or experiment; validate Marine/Phoenix 8a against established audio benchmarks.

The stratification protects against the two failure modes identified in the introduction. It prevents premature literalism by clearly separating Layer I (empirically established) from Layers III–IV (formal/philosophical). It prevents vacuous analogy by requiring, at Layer III, explicit mathematical content for each instantiation of $\mathcal{R}$.

Table 5: Extended concept atlas with epistemic layer.

Domain	**Admissibility detector**	**Failure mode**	**Layer**
GR / vacuum	Kretschmann scalar K	Curvature singularity	I / IV
λ-calculus	Type preservation $\Gamma \vdash M : T$	Non-termination	III
Sheaf semantics	$\check{H}^1 = 0$	Gluing obstruction	III
RSVP	$\dot{S} \leq \dot{\Sigma}_{\mathrm{crit}}$	Supercritical runaway	IV
Turbulence	Maximum dissipation rate	Sub-optimal α_b	I
Spherepop	Historical membrane coherence	Infinite bubbling	IV
MEM|8	Resonant retrieval boundedness	Attractor fragmentation	II
Hofstadter	Analogical coherence preservation	Recursive reminding	IV
Game theory	Iterated weak admissibility (IWD)	Dominated strategy	III
Observer (GR)	$g_{\mu\nu}V^\mu V^\nu > 0$	Trans-Planckian divergence	IV
Marine / Phoenix	Harmonic-lock stability	Lock failure / noise floor	II

15.1. The invariant that unifies all layers

Across all eleven domains in Table 5, the same structural claim recurs:

> **Stable systems preserve admissibility under local transformation.**

In λ-calculus, β-reduction preserves typing:

$$\Gamma \vdash (\lambda x.M)N : T \;\Rightarrow\; \Gamma \vdash M[x := N] : T.$$

In Spherepop, membrane collapse preserves historical topology. In Hofstadter, analogy preserves conceptual coherence under contextual rebinding. In RSVP, entropy gradients preserve viable future trajectories. In sheaf semantics, global sections exist only when $\check{H}^1(\mathcal{G}_\infty) = 0$. In quantum optics, interference preserves phase admissibility. In turbulence, maximum dissipation preserves the physically correct mixing trajectory. In game theory, iterated elimination preserves representation-independence of the solution. In relativistic observer theory, local Lorentz symmetry is preserved exactly within the admissible domain.

The admissibility principle is the claim that this invariant-preservation structure is not accidental but reflects a common operational grammar.

16. Admissibility and Constraint-Limited Observerhood: RSVP, CLIO, TARTAN, and Yarncrawler

Epistemic Status: Speculative Theoretical Proposal

This section integrates the admissibility program with several Flyxion frameworks (RSVP, CLIO, TARTAN, Yarncrawler) that are under active theoretical development. The connections drawn here are interpretive proposals, not established results.

16.1. Observerhood as emergent admissibility

The admissibility framework developed across the preceding sections aligns naturally with the RSVP program, the CLIO (*Constraint-Leveraged Inference Operator*) formalism, and the projection-based interpretation of cognition and physics developed throughout this work.

In RSVP, physical systems are not isolated objects moving through a passive geometric background. All observable structure emerges from constraint evolution within the scalar–vector entropy field $(\Phi, \mathbf{v}, S)$, where Φ denotes scalar coherence density, $\mathbf{v}$ directed constraint flow, and S admissible configurational entropy. Observerhood itself becomes a dynamically stabilized trajectory through constraint space. An observer is therefore not fundamental but emergent from persistent low-entropy coherence structures capable of recursively maintaining informational continuity.

The admissibility principle acquires a reformulation in RSVP language:

> An observer is physically admissible if and only if its trajectory remains recursively coherent under the local and global constraint geometry of the plenum.

The relativistic admissibility condition $g_{\mu\nu}V^\mu V^\nu > 0$ becomes only the lowest-order projection of a more general coherence requirement. The metric itself emerges as a coarse-grained geometric encoding of admissible informational trajectories within the plenum.

16.2. CLIO: cognition as constrained projection

In CLIO, cognition and measurement are sparse projection operators acting on inaccessible high-dimensional state spaces. Let

$$\pi : X \to M$$

denote the projection from a high-dimensional ontic trajectory manifold X into an operational measurement manifold M. The observer navigates compressed projections preserving only dynamically relevant constraint relations. Admissibility requires not merely geometric timelikeness but *projectional stability*:

$$\pi(R(x)) = \pi(x)$$

for admissible reduction operators R preserving coherent inferential structure.

This clarifies why Planck-scale divergences arise operationally rather than ontologically. Attempts to probe arbitrarily small scales force the observer to collapse the projection manifold faster than recursive coherence can be maintained. The failure is not a breakdown of spacetime itself but a breakdown of observer-stabilized projection consistency.

Entropy in this reading is not merely disorder but the logarithmic volume of dynamically accessible future trajectories:

$$S \sim \log |\mathcal{A}|,$$

where $\mathcal{A}$ is the admissible future trajectory set. Infinite boosts, singular localization, and trans-Planckian measurement attempts collapse $|\mathcal{A}|$ toward instability.

16.3. TARTAN and Yarncrawler

TARTAN (*Trajectory-Aware Recursive Tiling with Annotated Noise*) treats physical and cognitive systems as recursively tiled trajectory bundles carrying both dynamic evolution and semantic metadata. Observerhood corresponds to maintaining stable trajectory alignment across recursively compressed scales. Admissibility becomes equivalent to preserving trajectory continuity under multiscale coarse-graining.

Yarncrawler extends this to civilization-scale systems. Roads, infrastructures, memory systems, and institutions are trajectory stabilizers preserving admissible movement through physical and semantic space. A city ceases functioning not when matter disappears but when transport pathways lose coherent accessibility. Maintenance itself is a form of admissibility preservation: repair systems, garbage systems, and archival systems prevent collapse into inadmissible trajectory fragmentation.

16.4. The philosophical inversion

The synthesis yields a central philosophical inversion of the RSVP program:

> Physical law does not merely govern what exists; it governs what can remain coherently observable.

Admissibility is not an additional force or symmetry. It is a meta-constraint on coherent existence itself.

More precisely: the fundamental question of physics changes. Instead of asking only what mathematically exists, we begin asking what can remain coherently observable. The universe becomes not merely a collection of possible states but a dynamic filter selecting which trajectories can persist without collapsing their own conditions of observation.

This reformulation naturally integrates:

- relativistic geometry (admissibility as timelikeness and curvature boundedness);
- quantum operational limits (Planck scale as projectional collapse horizon);
- recursive cognition (thought as admissibility-maintenance under analogical pressure);
- informational thermodynamics (entropy as trajectory accessibility volume);
- social systems (institutions as admissibility-preserving trajectory stabilizers);
- semantic systems (meaning as globally glueable local consistency).

17. Discussion and Open Problems

17.1. What has been established

The following claims in this paper rest on established results:

- The λ-calculus formalism, β-reduction, Church–Rosser confluence, and Curry–Howard are standard (Church 1936; Sørensen–Urzyczyn 2006).
- The quantum dwell-time results (anomalous weak values, group-delay formulas) follow APL Quantum (2025).
- The DC-OPA performance figures follow the Nature Photonics report on the advanced BiBO/MgO:LN system.
- Hofstadter's theory of analogy and chunking is presented faithfully from published work (1979, 1995, 2013).
- The sheaf-theoretic gluing and obstruction framework is standard category theory (Kashiwara–Schapira 1994).

- The game-theoretic admissibility results (iterated weak dominance, representation-independence) follow Bicchieri–Schulte (1996).
- The turbulence validation results (maximum dissipation rate, $\alpha_b \approx 0.06$ matching RT experiment) follow Glimm et al. (preprint LA-UR-19-20285).
- The Kretschmann scalar as minimal non-degenerate curvature invariant is standard differential geometry.
- The multimodal lineage tracing methodology, computational frameworks (LineageOT, moslin, LineageVAE, PORCELAN), and characterization of progressive potency restriction follow Wang, He, and Hu (2026) in *Nature Reviews Genetics*. The interpretation of lineage tracing through the admissibility schema is the authors' proposal.
- The underconstraint argument follows Yudkowsky (2008); the endogenous growth and Moore's Law analyses are as described in that essay.

17.2. What is proposed but not yet established

- The Spherepop–λ-calculus formal equivalence requires a complete syntactic specification of Spherepop and a soundness proof for Table 1.
- The RSVP model parameters (λ_c, $\dot{\Sigma}_{\text{crit}}$) require empirical calibration or derivation from first principles.
- The Marine/Phoenix 8a architecture requires independent benchmark evaluation, ablation studies, and perceptual testing.
- The generalized confluence schema (Schema 1) is a conjecture; it holds in specific cases but has not been proved in general.
- The Hofstadter–Spherepop cognitive mapping is interpretive; it provides a framework for future empirical and formal work, not a proof.
- The observer-admissibility framework of Modgil–Patil (2026) requires covariant reconstruction of γ_g using tetrads or observer congruences; the global-bound formula is currently phenomenological.
- The Zenteno (2026) penalty-functional approach requires derivation of the penalty from an action principle.

17.3. Priority of the research program

Layer V of the five-layer architecture (Section 15) constitutes the forward research agenda. In order of logical priority:

1. Formalize Spherepop as an explicit substitution calculus and prove the correspondence with $\lambda\sigma$ or $\lambda\upsilon$.

2. Prove or refute the generalized confluence schema (Schema 1) for the RSVP relaxation operator under explicit stability assumptions.

3. Develop a covariant tetrad-based formulation of the generalized relativistic factor γ_g to put the observer-admissibility framework on rigorous footing.

4. Investigate whether the Kretschmann-scalar admissibility detector admits categorical generalization: what is the analogue of K for semantic, cognitive, and thermodynamic admissibility?

5. Apply the Bicchieri–Schulte recursion to multi-agent semantic systems; formalize whether LLM hallucination rates decrease under iterated admissibility elimination over agents' epistemic states.

6. Develop independent benchmarks for the Marine temporal-interval approach against STFT baselines on standardized audio corpora.

7. Formalize the sheaf-hallucination proposition and investigate testable predictions about failure modes in language models.

8. Explore whether Hofstadterian cognitive admissibility can be operationalized in a computational architecture (e.g., active inference or predictive coding frameworks).

17.4. Admissibility as research grammar rather than universal law

The reviewer of an early draft of this monograph offered the following observation: the paper's closing remark — that the admissibility principle is "a research grammar" rather than a physical law or mathematical theorem — is the paper's most accurate statement of its own contribution. That observation is correct, and the present work embraces rather than retreats from it.

The comparison to category theory is apt and worth dwelling on. Category theory did not claim that all mathematics was literally the same object, or that groups and topological spaces were formally equivalent. It supplied a *transport language* for structural relationships across domains: functors, natural transformations, adjoints, limits. This language proved extraordinarily productive not because it proved old theorems but because it made new questions visible.

The admissibility principle aims to occupy the same role. Its contribution is not a theorem but a vocabulary: admissibility region, perturbation, collapse operator, invariant, normal form, confluence, gluing obstruction. These terms, applied with the precision demanded by Definition 11.1 and the necessary conditions in the preceding subsection, allow researchers to ask coherent questions about constraint-preserving collapse in domains where those questions had not previously been posed in that form.

The most defensible version of the central thesis is therefore:

> Many systems — formal, physical, cognitive, social — exhibit recurring structural motifs involving constrained trajectory reduction, invariant preservation, and local compatibility propagation. These motifs are sufficiently precise to motivate a shared research grammar. Identifying the grammar is a contribution. Proving its instances are formally equivalent is the research program that follows from it.

18. Underconstrained Abstractions and the Limits of Historical Fit

18.1. The problem of historical underconstraint

One of the deepest epistemic problems facing any cross-domain synthesis is that historical data radically underdetermines abstraction. A single body of observations may admit many structurally compatible descriptions, and the mere existence of a retrospective fit does not establish that an abstraction captures the operative mechanism responsible for the observed behavior. A curve can be interpolated through finitely many points by polynomials of arbitrarily different character; the interpolation tells us nothing about which polynomial governs the generative process.

This problem was articulated with particular sharpness by Eliezer Yudkowsky in a 2008 essay titled "Underconstrained Abstractions," written as part of an extended exchange with Robin Hanson about the epistemic status of abstractions used in technological forecasting. Yudkowsky's central observation is worth quoting in its original form: "speaking of an abstraction being 'verified' by previous history is a tricky thing. There is this little problem of underconstraint — of there being more than one possible abstraction that the data 'verifies'." He continues: "The further away you get from highly regular things like atoms, and the closer you get to surface phenomena that are the final products of many moving parts, the more history underconstrains the abstractions that you use. This is part of what makes futurism difficult. If there were obviously only one story that fit the data, who would bother to use anything else?"

The problem is structural rather than merely practical. In domains close to fundamental physics, the regularity of the underlying processes tightly constrains the space of compatible abstractions. Atomic physics admits few retrospective stories because the entities involved obey highly symmetric and repeatable laws that sharply penalize deviant descriptions. But as one ascends toward emergent, historically layered systems — cognition, economics, institutional dynamics, semantic systems, civilization-scale technological development — the surface phenomena are the final product of many interacting layers, each of which introduces additional degrees of freedom that the historical record does not constrain. In such domains, retrospective fit is necessary but very far from sufficient. An abstraction that merely redescribes historical data without imposing constraints on what patterns it excludes is, in Yudkowsky's sense, underconstrained.

18.2. The endogenous growth illustration

Yudkowsky's primary illustration of the underconstraint problem involves the "endogenous growth" literature in economics, specifically theories in which knowledge accumulation is modeled as a combinatorial process. The seminal formulation treats ideas as being generated by combining other ideas N at a time, so that a stock of N available ideas produces on the order of $N!/(k!(N-k)!)$ potential new ideas for any fixed combination size k. From this combinatorial abundance, it follows that the number of available ideas will vastly exceed the number of people available to exploit them, and the rate of knowledge production will therefore be limited by the supply of researchers rather than by the supply of ideas.

Yudkowsky makes a precise and important observation about this structure. If the only proposition that actually enters the predictive machinery of the theory is that there exist more potential ideas than people to pursue them, then this is the only proposition that the theory's fit to historical data can even partially verify. The elaborate combinatorial apparatus — the N-choose-k framing, the specific functional form of the idea-generation process — is not doing any predictive work beyond what would be accomplished by the simpler statement: "I assume there are more ideas than people to exploit them." Any hypothesis with "more ideas than people to exploit them" in its structure will fit the same historical data equally well. The additional mathematical specificity is not constrained by the data because it is not used in the forward application; it provides a causal-sounding story without providing causal leverage.

This is the core epistemic failure Yudkowsky identifies: an abstraction can achieve the appearance of precision and the appearance of historical verification while remaining genuinely underconstrained, because the constrained parts of the abstraction are not the parts doing the predictive work. The decorative mathematics acquires false credibility by association with the parts of the theory that are actually tested.

18.3. Moore's Law and multiple compatible abstractions

The Moore's Law case is the most concrete illustration in Yudkowsky's essay, and it bears careful examination because it directly concerns the type of technological forecasting that has the highest stakes for arguments about artificial intelligence.

The historical trajectory of exponential growth in transistor density is compatible with at least four distinct generating abstractions, all fitting the same observed data. One abstraction treats transistor count as a function of elapsed calendar time, an empirical regularity with no identified causal mechanism beyond the observation itself. A second treats it as a consequence of exponential growth in capital investment in semiconductor manufacturing and research and development, predicting continuation as long as that investment continues. A third treats it as a function of optimization pressure applied by a community of human researchers operating at a fixed cognitive rate, predicting continuation as long as the research community remains adequately populated and funded. A fourth treats it as involving some degree of recursive amplification, in which improved tools make further improvement easier, predicting a non-linear relationship between the trajectory of improvement and the substrate on which the researchers themselves operate.

Yudkowsky raises the counterfactual that makes these abstractions empirically distinguishable in principle while remaining indistinguishable within the historical record: "Would the further progress of Moore's Law have been different from that in our own world, relative to sidereal time, if, for these last fifty years, the researchers themselves had been running on the latest generation of computer chip at any given point?" This counterfactual is not answerable from historical data, because the researchers were not in fact running on progressively faster hardware during the period when the historical data was generated. But the four competing abstractions give sharply different answers to it, and those different answers imply radically different predictions about what would happen if the researchers' cognitive substrate were to change.

The underconstrained-abstractions problem is therefore not merely that multiple stories fit the data. It is that the stories diverge precisely where they matter most for forward application, and the historical data provides no basis for choosing between them in that region.

18.4. Moravec and the forward-application problem

Yudkowsky's essay also references Moravec's 1988 work as an illustration of a complementary failure: the problem of assuming a unique forward application of an abstraction that has, in fact, multiple possible extensions. Moravec estimated the computing power required for human-level artificial intelligence by multiplying the apparent number of operations performed by the human brain per second by an appropriate efficiency factor, arriving at a figure he then extrapolated forward using Moore's Law to estimate the arrival date. Yudkowsky notes that Moravec "spent a lot of time talking about how much 'computing power' the human brain seems to use — but much less time talking about whether an AI would use the same amount of computing power, or whether using Moore's Law to extrapolate the first supercomputer of this size is the right way to time the arrival of AI."

The problem is not that Moravec's historical estimates of brain computational equivalence were wrong, but that the forward application contained a hidden assumption: that an artificial system performing at human level would require approximately the same number of operations per second as the biological system being emulated, and that the arrival of sufficient hardware therefore determines the arrival of the capability. This assumption is not constrained by the historical data used to establish the computational-equivalence estimate. It is an additional hypothesis that enters silently at the point of forward application, and it happens to be the hypothesis doing all the predictive work in the arrival-date calculation. The estimate is therefore not a prediction so much as a restatement of the hidden assumption, dressed in the credibility borrowed from the carefully validated computational-equivalence estimate.

This is what Yudkowsky calls the problem of "a lot of slack with regards to exactly what will happen with respect to that abstraction, going forward": after constructing a historical story that appears to verify an abstraction, the forward application typically requires additional choices about how to apply the abstraction, and those choices are not themselves constrained by the historical verification. The appearance of a single determinate prediction conceals an underspecified space of possible applications.

18.5. The admissibility response to underconstraint

The admissibility framework attempts to respond to this problem not by eliminating abstraction but by constraining which kinds of abstractions count as legitimate transport structures.

An admissible abstraction must do more than fit historical data. It must preserve identifiable invariants under perturbation, specify failure modes precisely enough that they could in principle be observed, restrict future trajectories in ways that generate differential predictions across competing causal stories, and produce exclusions rather than merely redescriptions. An abstraction that imposes no exclusions is vacuous; it is capable of absorbing any observation after the fact by adjusting its parameters or reinterpreting its terms. The formal apparatus of Definition 11.1 and the augmented strong admissibility criterion of Appendix F are therefore not merely mathematical decoration. They are the mechanisms by which the framework resists the slide into explanatory vacuity.

The connection to Yudkowsky's diagnosis is direct. An underconstrained abstraction in his sense is precisely an abstraction that fails to specify which of its components are doing the predictive work and which are providing a causal-sounding story without causal leverage. The admissibility framework operationalizes this distinction by requiring that the collapse operator, the compatibility condition, and the invariant be made explicit and distinct. An abstraction in which the "invariant" is never operationalized, in which the "compatibility condition" is satisfied by any input, or in which the "collapse operator" is never applied to generate a prediction that could fail, is an abstraction that fails the admissibility criterion. It may fit the history, but it is not doing so through a mechanism that constrains the future.

18.6. The danger of unconstrained universality for the present framework

This diagnosis applies reflexively to the admissibility principle itself. If the admissibility framework were to become underconstrained in Yudkowsky's sense, it would fit all observed historical instances of "constraint-preserving transformation" while providing no leverage on which future processes will or will not exhibit admissibility structure. The framework would then be exactly the kind of abstraction it is designed to discipline.

The non-vacuity criterion of Appendix F and the three augmented conditions of strong admissibility are the framework's attempt to avoid this failure. Requiring that the admissibility predicate recursively restructure the future trajectory space, that the invariant be transportable and usable by downstream processes, and that the system exhibit genuine path-sensitivity are requirements specifically designed to ensure that the admissibility framework is not merely a retrospective vocabulary. A thermostat satisfies a liberal reading of Definition 11.1 but fails strong admissibility precisely because no discriminatory work is being done: the thermostat's future behavior is fully determined by its instantaneous state, there is no invariant being transported downstream, and there is no path-sensitivity. Excluding the thermostat is not an embarrassment; it is the framework demonstrating that it imposes genuine constraints.

The same logic applies to the cross-domain cases treated throughout this monograph. The sheaf-theoretic hallucination model, the Glimm turbulence admissibility criterion, the Bicchieri–Schulte recursive epistemic admissibility, and the Kretschmann scalar curvature detector are all cases where the admissibility structure is specified with enough precision to be wrong. The RSVP framework and Spherepop are cases where the specification is not yet sufficiently precise, and the monograph identifies them explicitly as research programs rather than established results. This stratification is the monograph's attempt to practice what it preaches: not all cross-domain analogies are equally constrained, and the framework's own epistemic standards require acknowledging the difference.

18.7. Forecasting as constraint-topology tracking

The future is difficult to predict not merely because it has not yet occurred but because the admissibility manifold may itself change. Historical regularities emerge from fixed constraints, contingent optimization pressures, hidden mesoscale structures, or temporary equilibrium conditions, and when any of these underlying structures shifts, an abstraction previously fitted to the historical trajectory may cease to apply not gradually but catastrophically.

This is why extrapolation across phase transitions in the admissibility manifold is unreliable even for abstractions with excellent retrospective fit. A system near a bifurcation point in its constraint topology generates historical data that looks regular and predictable, because the constraint structure has been stable throughout the observed period. But an abstraction fitted to that period carries no information about the system's behavior after the bifurcation. The admissibility principle therefore treats forecasting as a problem of tracking constraint evolution rather than extending historical curves. Yudkowsky's question about Moore's Law in the counterfactual world where researchers run on their own products is precisely the question of whether the constraint topology would have changed, and if so, how the abstraction would need to be modified to remain valid under the new topology. Answering it requires understanding what generates the observed regularities, not merely that they have been observed.

18.8. Connection to sheaf-theoretic obstruction

The problem of underconstrained abstraction is closely related to the sheaf-theoretic notion of gluing failure developed in Section 11 and Appendix C.

An abstraction may remain locally compatible with historical data while failing to extend coherently across perturbative or counterfactual domains. The historical record constitutes a collection of local sections $\{s_i\}$ over an open cover of the observation domain, and the abstraction proposes that these sections glue into a global section extending beyond the observed region. The existence of pairwise-compatible local sections does not guarantee the existence of a global section: the obstruction to extension is measured by the first Čech cohomology class $[\omega] \in \check{H}^1(X, \mathcal{F})$.

When $[\omega] \neq 0$, the abstraction glues coherently over the observed history but fails globally once the constraint topology changes. This is the sheaf-theoretic reading of Yudkowsky's underconstraint problem: the abstraction satisfies the cocycle condition on observed overlaps but cannot be extended to a globally coherent section of the semantic manifold. The forward-application failures Yudkowsky identifies — Moravec's computational extrapolation, endogenous growth models applied to recursive self-improvement, Moore's Law applied across a change in the cognitive substrate of researchers — are all cases where local historical compatibility does not imply global extension. The obstruction is invisible within the range of observations used to construct the abstraction and becomes apparent only in the counterfactual or future domain where the abstraction is actually needed.

The admissibility framework offers a partial response. By requiring that an admissible abstraction specify its failure modes explicitly, it forces the question of where the global section might fail to extend. An abstraction that specifies the conditions under which its invariant ceases to be preserved, the boundary of its admissibility region, and the perturbations that fall outside its compatibility condition is an abstraction whose obstruction classes are at least partially identified before the failure occurs. This does not eliminate the problem of underdetermination, but it converts a silent failure mode into a visible research question: whether the abstraction extends globally becomes a question with specific mathematical content rather than an implicit assumption whose failure is discovered only after the extrapolation has been committed to.

18.9. Closing remark

The admissibility principle is not a physical law, a mathematical theorem, or an empirical discovery. It is a *research grammar*: a way of asking the same question across very different systems. That question is: *what makes a perturbation locally receivable, and what becomes stable once it has been received?*

When that question can be answered precisely for a given domain, the admissibility schema becomes a theorem. When it can be answered approximately, it becomes a model. When the answer remains unclear, it remains heuristic.

The purpose of this monograph is to make the question precise enough that the distinction between these three cases can be drawn.

References

- Church, A. (1936). An unsolvable problem of elementary number theory. *American Journal of Mathematics*, 58(2), 345–363.
- Curry, H. B., & Feys, R. (1958). *Combinatory Logic*, Vol. I. North-Holland.
- Girard, J.-Y., Taylor, P., & Lafont, Y. (1989). *Proofs and Types.* Cambridge University Press.
- Hofstadter, D. R. (1979). *Gödel, Escher, Bach: An Eternal Golden Braid.* Basic Books.
- Hofstadter, D. R., & Mitchell, M. (1995). *Fluid Concepts and Creative Analogies.* Basic Books.
- Hofstadter, D. R., & Sander, E. (2013). *Surfaces and Essences: Analogy as the Fuel and Fire of Thinking.* Basic Books.
- Howard, W. A. (1980). The formulae-as-types notion of construction. In J. P. Seldin & J. R. Hindley (Eds.), *To H. B. Curry: Essays on Combinatory Logic, Lambda Calculus, and Formalism* (pp. 479–490). Academic Press.
- Sørensen, M. H., & Urzyczyn, P. (2006). *Lectures on the Curry–Howard Isomorphism.* Elsevier (Studies in Logic and the Foundations of Mathematics, Vol. 149).

APL Quantum Author(s). (2025). Atomic excitation times for transmitted and scattered photons. *APL Quantum*, **2**, 036108. [Full author list to be inserted.]

Photonics DC-OPA Author(s). (2024–25). Advanced dual-chirped optical parametric amplification producing single-cycle MIR pulses. *Nature Photonics.* [Full citation to be inserted.]

- Kashiwara, M., & Schapira, P. (1994). *Sheaves on Manifolds.* Springer.
- de Bruijn, N. G. (1972). Lambda calculus notation with nameless dummies. *Indagationes Mathematicae*, 34(5), 381–392.
- Abadi, M., Cardelli, L., Curien, P.-L., & Lévy, J.-J. (1991). Explicit substitutions. *Journal of Functional Programming*, 1(4), 375–416.
- Bicchieri, C., & Schulte, O. (1996). Common reasoning about admissibility. *Erkenntnis* (revised 1996, forthcoming).
- Evangelista, R. (preprint). Admissibility and the minimal curvature bound in covariant theories. [Author affiliation not given; full citation to be completed.]

- Glimm, J., Cheng, B., Sharp, D. H., & Kaman, T. A crisis for the verification and validation of turbulence simulations. Los Alamos National Laboratory Preprint LA-UR-19-20285. [Published version citation to be inserted.]
- Johansson, I. (year). Construction as reduction. [Full citation to be inserted — referenced for the analogy-as-reduction framing in linguistics.]
- Hawking, S. W. (1975). Particle creation by black holes. *Communications in Mathematical Physics*, 43, 199–220.
- Bekenstein, J. D. (1973). Black holes and entropy. *Physical Review D*, 7, 2333–2346.
- Bousso, R. (2002). The holographic principle. *Reviews of Modern Physics*, 74, 825–874.
- Jacobson, T. (1995). Thermodynamics of spacetime: the Einstein equation of state. *Physical Review Letters*, 75, 1260–1263.
- de Lellis, C., & Székelyhidi, L. (2009). The Euler equations as a differential inclusion. *Annals of Mathematics*, 170, 1417–1436. [Cited for wild Euler-equation solutions establishing nonuniqueness; background for the Glimm et al. turbulence admissibility argument.]
- Modgil, M. S., & Patil, D. D. (2026). On the invariance of Planck scale under special relativity and the geometry of quantum space. Cosmos Research Labs. [Preprint; February 9, 2026.]
- Paton, A. J. (2026). Structural admissibility as a pre-explanatory condition in scientific reasoning. The Paton System.
- Wang, K., He, X., & Hu, Z. (2026). Computational approaches for multimodal lineage tracing. *Nature Reviews Genetics.* 10.1038/s41576-026-00969-9 [Primary source for the biological admissibility instantiation in Section 6.]
- Yudkowsky, E. (2008). Underconstrained Abstractions. *LessWrong* (December 4, 2008). `https://www.lesswrong.com/posts/underconstrained-abstractions` [Cited for the problem of historical underconstraint in technological forecasting, the endogenous growth illustration, and the Moore's Law multiple-abstraction analysis.]
- Moravec, H. (1988). *Mind Children: The Future of Robot and Human Intelligence.* Harvard University Press. [Cited for the computational-equivalence extrapolation and forward-application problem discussed in Section 18.]

- Zenteno, C. (2026). Admissibility under constraint: dynamical criterion beyond existence and consistency. Zoa Industries — Compressed Consciousness Laboratory. Zenodo. `https://doi.org/10.5281/zenodo.18370414`

A. Category-Theoretic Formulation of Admissibility

This appendix develops a category-theoretic formulation of the generalized reduction operator introduced in Section 11. The goal is not to assert a completed categorical foundation for the admissibility principle, but to identify the minimal mathematical structures such a foundation would require, and to demonstrate that the framework is at least as rich as established categorical machinery in several limiting cases.

A.1. Admissibility categories

Let $\mathcal{C}$ be a category whose objects are admissibility regions $\mathcal{B}$ and whose morphisms represent admissibility-preserving transformations between them. We equip $\mathcal{C}$ with three additional pieces of structure.

The first is a partial reduction endofunctor

$$\mathfrak{R} : \mathcal{C} \rightharpoonup \mathcal{C},$$

defined only on those objects $\mathcal{B}_i$ for which there exists an admissible perturbation Δ_i compatible with the constraint structure $\partial(\mathcal{B}_i)$. Partiality here is essential: the functor is not globally defined, and the domain of $\mathfrak{R}$ is determined by the admissibility predicate rather than by universal coverage.

The second piece of structure is a family of admissibility predicates $\mathsf{Adm}_{A,B}(f)$ on morphisms $f : A \to B$, specifying which morphisms participate in coherent collapse sequences. A morphism f is admissible if the image $\mathfrak{R}(A)$ is reachable from A via f without violating the invariant condition.

The third is an invariant-preservation natural transformation

$$\eta : \iota \Rightarrow \iota \circ \mathfrak{R},$$

where $\iota : \mathcal{C} \to \mathcal{I}$ is a functor into an invariant category $\mathcal{I}$ tracking the transportable structure that collapse must preserve. The naturality of η encodes the requirement that invariant preservation is uniform across all admissible reductions: no reduction may preserve the invariant selectively for some morphisms while violating it for others.

Definition A.1 (Admissibility category)**.** An *admissibility category* is a tuple $(\mathcal{C}, \mathfrak{R}, \mathsf{Adm}, \iota, \eta)$ in which $\mathcal{C}$ is a category, $\mathfrak{R}$ is a partial endofunctor, Adm is a family of admissibility predicates on morphisms, ι is an invariant functor, and η is a natural transformation $\eta_A : \iota(A) \to \iota(\mathfrak{R}(A))$ for every A in the domain of $\mathfrak{R}$, with each η_A an isomorphism in $\mathcal{I}$.

The isomorphism condition on η_A captures the idea that admissible reduction does not merely approximate the invariant but preserves it exactly. In typed λ-calculus, this corresponds to the Subject Reduction theorem: β-reduction preserves typing judgements, so the type is the invariant and the preservation is exact under the isomorphism $\eta_A : T_A \xrightarrow{\sim} T_{A'}$. In sheaf semantics it corresponds to the requirement that gluing maps are isomorphisms on overlapping sections.

A.2. Ordinary λ-calculus as a special case

The untyped λ-calculus corresponds to the special case in which the admissibility predicate is trivially satisfied after α-conversion: every redex is admissible, $\mathfrak{R}$ is defined everywhere on reducts, and the invariant category $\mathcal{I}$ is trivial, reflecting the fact that untyped reduction carries no non-trivial invariant beyond the reduction relation itself.

In the simply-typed setting, $\mathcal{I}$ becomes the category of simple types and type-preserving substitutions, and Subject Reduction provides η. In System F, $\mathcal{I}$ becomes the category of polymorphic types. In dependent type theory, $\mathcal{I}$ becomes a category whose objects are dependent types and whose morphisms are definitional equalities. The admissibility framework thus interpolates cleanly across the standard hierarchy of type theories, interpreting each as an admissibility category with progressively richer invariant structure.

Spherepop generalizes this by making admissibility depend on historical topology. The object $A = (\mathcal{I}(A), \partial(A), H(A))$ in an admissibility category for Spherepop carries a third component, the history $H(A)$, which enters the admissibility predicate: $\mathsf{Adm}_{A,B}(f)$ holds only if f is consistent with $H(A)$ as a causal continuation. This is the categorical expression of condition (iii) in Definition 11.1.

A.3. Confluence as a weak diamond condition

The Church–Rosser property admits a clean categorical reformulation. In the admissibility category, confluence becomes the following.

Definition A.2 (Weak admissibility diamond)**.** An admissibility category $(\mathcal{C}, \mathfrak{R}, \mathsf{Adm}, \iota, \eta)$ satisfies the *weak admissibility diamond property* if for every pair of admissible morphisms $f : A \to B$ and $g : A \to C$ in $\mathcal{C}$, there exists an object D together with admissible morphisms $f' : B \to D$ and $g' : C \to D$ such that $f' \circ f = g' \circ g$ holds up to the equivalence relation generated by the admissibility predicates.

This definition generalizes Church–Rosser confluence, sheaf gluing conditions, coherence diagrams in monoidal categories, and path reconciliation in homotopy theory. In each case the weak admissibility diamond asserts that locally incompatible collapse paths can be joined into a common successor. These are not the same theorem: each domain proves the diamond by entirely different arguments. The admissibility category framework supplies a common *language* in which the analogous structural condition in each domain is stated in the same syntactic form.

A.4. Functorial transport between admissibility categories

A central aspiration of the monograph is that admissibility structure can be transported between domains. In categorical terms this requires functors $F : \mathcal{C}_1 \to \mathcal{C}_2$ between admissibility categories that intertwine the respective reduction functors and preserve the admissibility predicates.

A functor F between admissibility categories is *admissibility-preserving* if there exists a natural transformation

$$\phi : F \circ \mathfrak{R}_1 \Rightarrow \mathfrak{R}_2 \circ F$$

making the reduction square commute, and if F maps admissible morphisms in $\mathcal{C}_1$ to admissible morphisms in $\mathcal{C}_2$.

The transport claims made in the main text correspond to the hypothesis that admissibility-preserving functors exist between the admissibility categories determined by each domain. Verifying this hypothesis in any specific case is a well-posed mathematical problem and constitutes a research target within the Layer V program.

B. Fiber Bundles, Projection, and Semantic Collapse

This appendix develops the fiber-bundle interpretation of admissibility collapse, giving a geometric formulation of the distinction between admissibility-preserving compression and destructive projection.

B.1. The projection setup

Let $\pi : X \to M$ be a smooth surjection from a high-dimensional trajectory space X onto a lower-dimensional manifold M. Interpret X as encoding the full historical microstructure of a system, while M encodes the compressed observable structure accessible to an external observer or a downstream process. For each point $m \in M$, the fiber $\pi^{-1}(m)$ consists of all microscopic states compatible with macroscopic representation m.

A section $s : M \to X$ assigns to each macroscopic state a canonical microscopic representative. The existence of globally consistent sections is the analog in this geometric language of the existence of global sections in sheaf theory. When no global section exists, locally consistent microstructure cannot be assembled into a globally coherent representation.

Definition B.1 (Admissible projection)**.** The projection $\pi : X \to M$ is *admissible* if the following conditions hold. First, admissible reductions in X project to continuous paths in M: if $\gamma : [0,1] \to X$ is an admissible trajectory then $\pi \circ \gamma$ is a continuous path in M. Second, admissible reductions contract fibers monotonically: if $x_0, x_1 \in X$ are connected by an admissible trajectory and $\pi(x_0) = \pi(x_1) = m$, then the fiber $\pi^{-1}(m)$ does not expand under the reduction. Third, invariant structure is transportable across local trivializations: for every local trivialization $\phi_i : \pi^{-1}(U_i) \to U_i \times F$ of the bundle, the invariant functor ι commutes with the transition maps $\phi_j \circ \phi_i^{-1}$ on overlapping charts $U_i \cap U_j$.

The third condition is precisely the cocycle condition for a fiber bundle with structure group determined by the admissibility predicates. Its failure corresponds to the gluing obstruction: locally admissible structure that cannot be assembled coherently.

B.2. Curvature of the admissibility bundle

When the admissibility bundle $\pi : X \to M$ admits a connection ∇, its curvature two-form $\Omega \in \Omega^2(M, \mathrm{End}(F))$ measures the failure of parallel transport to be path-independent. Two paths γ_1, γ_2 in M from m_0 to m_1 produce holonomy operators $P_{\gamma_1}, P_{\gamma_2}$ on the fiber $\pi^{-1}(m_0)$; their difference is encoded by the holonomy group and controlled by Ω.

In the admissibility framework, non-vanishing curvature on admissible paths corresponds to path-dependence in admissibility transport: two reduction sequences beginning at the same admissibility region may reach structurally inequivalent successors even if both are individually admissible. Confluence is therefore the flatness condition on the admissibility connection restricted to admissible paths: the curvature form Ω vanishes on pairs of tangent vectors determined by admissible reductions.

B.3. Three failure modes as geometric pathologies

Three distinct failure modes become precise in this geometric language.

Projection collapse occurs when distinct states $x_0, x_1 \in X$ satisfy $\pi(x_0) = \pi(x_1)$ even though $x_0 \neq x_1$. Their admissibility distinction is erased by the projection, and subsequent evolution over M cannot distinguish consequences that were distinguishable in X. This is the geometric form of lossy compression: the map discards information that matters for future admissibility.

Fiber fragmentation occurs when the fiber $\pi^{-1}(m)$ splits into disconnected components under reduction, so that globally consistent microscopic representatives become unavailable even though local sections still exist. The system loses track of the causal relationships between its own microstructures.

Topological hallucination occurs when local sections $s_i : U_i \to \pi^{-1}(U_i)$ exist for every open set in an open cover of M and agree pairwise on overlaps $U_i \cap U_j$, but no global section $s : M \to X$ exists. The first Čech cohomology $\check{H}^1(M, \mathcal{F})$ is non-zero, where $\mathcal{F}$ is the sheaf of local sections. This is the geometric characterization of hallucination as a topological obstruction rather than as a local consistency failure.

B.4. The map-territory instability as bundle degeneration

When a system begins optimizing over M while treating M as identical to X, the bundle structure degenerates. The admissibility predicates defined relative to X become redefined relative to M, the connection on the bundle is replaced by a degenerate flat connection, and the transition maps $\phi_j \circ \phi_i^{-1}$ become trivial, collapsing the non-trivial fiber geometry into a product bundle and erasing the admissibility-bearing structure of the fibers.

This is the geometric form of Goodhart's Law: optimizing a measure destroys its correlation with the underlying structure it was intended to represent. The admissibility framework implies that any representation subsequently subjected to optimization pressure must preserve non-trivial fiber geometry from the outset. Representations that are flat by construction cannot support coherent admissibility transport under optimization.

C. Sheaf Cohomology, Semantic Coherence, and Hallucination

This appendix develops the sheaf-cohomology interpretation of semantic coherence and hallucination in detail, extending the Proposition in Section 11 to a full cohomological analysis.

C.1. Presheaves and sheaves on semantic manifolds

Let X be a topological space representing a semantic manifold and let $\mathcal{U} = \{U_i\}_{i \in I}$ be an open cover of X. Let $\mathcal{F}$ be a presheaf on X with values in the category of semantic structures, so that $\mathcal{F}(U)$ consists of locally consistent semantic assignments over U and the restriction maps $\rho_{UV} : \mathcal{F}(U) \to \mathcal{F}(V)$ for $V \subseteq U$ encode semantic specialization.

The presheaf $\mathcal{F}$ is a sheaf if two conditions hold. The locality condition asserts that sections agreeing on every element of a cover are equal: if $s, t \in \mathcal{F}(U)$ with $\rho_{U,U_i}(s) = \rho_{U,U_i}(t)$ for all i, then $s = t$. The gluing condition asserts that a compatible family $\{s_i \in \mathcal{F}(U_i)\}$ satisfying $\rho_{U_i,U_i \cap U_j}(s_i) = \rho_{U_j,U_i \cap U_j}(s_j)$ for all i, j assembles into a unique global section $s \in \mathcal{F}(X)$ with $\rho_{X,U_i}(s) = s_i$.

C.2. Čech cohomology and obstruction theory

For an open cover $\mathcal{U} = \{U_i\}_{i \in I}$, the Čech cochain groups are

$$\check{C}^n(\mathcal{U}, \mathcal{F}) = \prod_{i_0 < i_1 < \cdots < i_n} \mathcal{F}(U_{i_0} \cap U_{i_1} \cap \cdots \cap U_{i_n}),$$

with coboundary maps $\delta^n : \check{C}^n \to \check{C}^{n+1}$ defined by alternating restriction maps in the standard way. The cohomology groups are

$$\check{H}^n(\mathcal{U}, \mathcal{F}) = \ker \delta^n / \operatorname{im} \delta^{n-1}.$$

The zeroth group $\check{H}^0(X, \mathcal{F})$ is the group of global sections. The first group $\check{H}^1(X, \mathcal{F})$ measures the obstruction to assembling locally coherent patches into a global section: a one-cocycle $\{s_{ij} \in \mathcal{F}(U_i \cap U_j)\}$ satisfying

$$s_{jk} - s_{ik} + s_{ij} = 0 \in \mathcal{F}(U_i \cap U_j \cap U_k)$$

is a coboundary if and only if there exist $\{t_i \in \mathcal{F}(U_i)\}$ with $s_{ij} = t_j - t_i$. Non-coboundary cocycles are the obstructions.

Definition C.1 (Semantic hallucination)**.** A *semantic hallucination* with respect to the cover $\mathcal{U}$ and presheaf $\mathcal{F}$ is a non-zero element of $\check{H}^1(X, \mathcal{F})$: an isomorphism class of Čech one-cocycles that cannot be written as coboundaries. Equivalently, it is a collection of locally compatible semantic patches that cannot be assembled into a globally coherent semantic structure.

The characteristic phenomenology of hallucination follows directly. Each local patch $s_i \in \mathcal{F}(U_i)$ is internally consistent. The pairwise compatibility conditions are satisfied on every overlap. But the global section fails to exist because the transition data form a non-trivial cohomology class. Local inspection cannot detect the obstruction; only a global coherence check across triple overlaps can reveal it.

C.3. Admissibility as exactness and higher obstruction

The vanishing of $\check{H}^1(X, \mathcal{F})$ is equivalent to exactness of the sequence

$$0 \to \check{H}^0(X, \mathcal{F}) \to \prod_i \mathcal{F}(U_i) \xrightarrow{\delta^0} \prod_{i<j} \mathcal{F}(U_i \cap U_j).$$

This exactness condition is the sheaf-theoretic formulation of semantic admissibility: every locally admissible structure passing the compatibility test at overlaps must be globally admissible. The admissibility framework asserts that a semantic system is admissibility-sound if and only if $\check{H}^1(X, \mathcal{F}) = 0$.

The higher groups $\check{H}^n$ for $n \geq 2$ measure cascaded obstructions. A non-trivial class in $\check{H}^2$ indicates failures of compatibility among triple overlaps assembled from locally consistent pairwise data, corresponding in language models to systems producing individually plausible facts that are pairwise consistent but collectively contradictory through their logical implications.

D. Generalized Entropy-Constraint Dynamics

This appendix develops a partial mathematical structure for RSVP-style entropy-constraint dynamics, making the admissibility functional and variational equations as precise as the current state of the framework permits.

D.1. The RSVP field system

Let $\Omega \subseteq \mathbb{R}^n$ be a compact region with smooth boundary, representing the spatial domain of the plenum. The state at time t is described by the scalar coherence potential $\Phi : \Omega \times [0,T] \to \mathbb{R}$, the vector flow field $\mathbf{v} : \Omega \times [0,T] \to \mathbb{R}^n$, and the entropy density $S : \Omega \times [0,T] \to \mathbb{R}_{\geq 0}$.

The *admissibility functional* is defined as

$$\mathcal{A}[\Phi, \mathbf{v}, S] = \int_\Omega \left(\alpha\, |\nabla\Phi|^2 + \beta\, |\nabla \times \mathbf{v}|^2 + \gamma\, S + \mu\, (\nabla \cdot \mathbf{v})^2\right) dV,$$

where $\alpha, \beta, \gamma, \mu > 0$ are coupling constants. The four terms penalize respectively: scalar gradient concentration, vortical flow structure, entropy density, and compressional modes in $\mathbf{v}$.

Admissible trajectories satisfy two simultaneous constraints. The variational inequality $\frac{d}{dt}\mathcal{A}[\Phi(t), \mathbf{v}(t), S(t)] \leq 0$ asserts that the admissibility functional decreases along the trajectory. The local entropy production constraint

$$\dot{S}(x,t) = \nabla \cdot \mathbf{J}_S(x,t) + \sigma(x,t) \geq 0,$$

combined with the global stability condition

$$\int_\Omega \sigma(x,t)\, dV < \dot{\Sigma}_{\text{crit}},$$

prevents supercritical entropy production.

D.2. Euler–Lagrange equations for admissible attractors

Setting the first variation of $\mathcal{A}$ to zero yields the Euler–Lagrange system. Varying with respect to Φ gives

$$-2\alpha\, \Delta\Phi = \lambda_\Phi,$$

where λ_Φ is a Lagrange multiplier enforcing the entropy coupling. Varying with respect to $\mathbf{v}$ gives, for divergence-free flow fields,

$$2\beta\,\Delta\mathbf{v} = \boldsymbol{\lambda}_v.$$

An *admissible attractor* $(\Phi^*, \mathbf{v}^*, S^*)$ satisfies these equations with $\lambda_\Phi = \gamma$ and $\boldsymbol{\lambda}_v = \mathbf{0}$:

$$\Delta\Phi^* = -\frac{\gamma}{2\alpha}, \qquad \Delta\mathbf{v}^* = \mathbf{0},$$

together with $\int_\Omega \sigma^*\,dV < \dot{\Sigma}_{\mathrm{crit}}$.

D.3. Stability analysis and the critical coupling

To analyze stability of $(\Phi^*, \mathbf{v}^*, S^*)$ under perturbations $(\delta\Phi, \delta\mathbf{v}, \delta S)$, expand $\mathcal{A}$ to second order:

$$\delta^2\mathcal{A} = \int_\Omega \big(\alpha\,|\nabla(\delta\Phi)|^2 + \beta\,|\nabla\times(\delta\mathbf{v})|^2 + \gamma\,\delta S\big)\,dV.$$

Since $\alpha, \beta > 0$, the first two terms are positive-definite. The stability condition reduces to $\gamma\int_\Omega \delta S\,dV > 0$ for admissible entropy perturbations. Since admissibility requires $\delta S \geq 0$ pointwise, the condition is equivalent to $\gamma > 0$, which holds by assumption. Every admissible attractor is therefore a local minimum of $\mathcal{A}$ within the admissible sector.

The parameter λ_c of the main text can be identified with

$$\lambda_c = \frac{\gamma}{2\alpha},$$

which controls the magnitude of the uniform Laplacian of Φ^* and sets the characteristic length scale of the coherence field. For $\lambda_c \approx 0.42$, the coherence field has a characteristic wavelength consistent with stable attractor formation in the model. This is the sense in which λ_c is a heuristic scaling parameter: it encodes the geometry of admissible attractors within the current variational framework, not an empirical universal constant.

E. Homotopy-Theoretic Reading of Identity Persistence

This appendix develops the homotopy-theoretic interpretation of identity under transformation, arguing that the admissibility framework requires replacing extensional identity with homotopical identity.

E.1. The inadequacy of extensional identity

Classical symbolic systems identify objects by their extensional properties: two objects are the same if they have the same parts or satisfy the same predicates. This identification is static and ahistorical. The admissibility framework requires a different notion because condition (iii) of Definition 11.1 asserts that history is strictly extended by each reduction: the same extensional state reached by two different admissibility trajectories may carry different structural consequences.

Formally, let $x_0, x_1 \in X$ be two points in the trajectory space with $\pi(x_0) = \pi(x_1) = m \in M$, so that they are extensionally identical at the level of the compressed representation. The admissibility framework distinguishes them if the homotopy classes of their approach trajectories differ: the path taken to reach x_0 and the path taken to reach x_1 may have deposited different admissibility history, making the two points structurally inequivalent despite extensional coincidence.

Definition E.1 (Admissible identity path)**.** An *admissible identity path* from x_0 to x_1 is a continuous map $\gamma : [0,1] \to X$ with $\gamma(0) = x_0$ and $\gamma(1) = x_1$ such that every intermediate state $\gamma(t)$ is admissible for all $t \in [0,1]$.

Two states are *admissibly homotopic* if there exists an admissible identity path between them. The admissibly homotopic equivalence relation partitions X into connected components of the admissible locus X_{adm}.

E.2. The fundamental group of the admissible locus

The fundamental group $\pi_1(X_{\mathsf{adm}}, x_0)$ encodes homotopy classes of admissible loops based at x_0. In the Spherepop interpretation, a non-trivial loop corresponds to a sequence of membrane collapses that returns the bubble system to its original topology but not its original state: the history has changed even though the extensional description has not.

In the cognitive interpretation, a non-trivial loop in the admissible conceptual space corresponds to a thought trajectory returning to an apparently familiar concept while arriving from a different direction, encoding a different set of analogical associations. Hofstadter's strange loops are precisely the non-trivial elements of $\pi_1(X_{\mathsf{adm}}, x_0)$.

E.3. Higher homotopy groups and coherence conditions

The higher homotopy groups $\pi_n(X_{\mathsf{adm}}, x_0)$ for $n \geq 2$ encode higher-dimensional coherence conditions. Vanishing of π_2 asserts that any two admissible homotopies between the same pair of paths can themselves be admissibly homotoped to each other, a condition analogous to Church–Rosser: just as Church–Rosser asserts that reduction sequences from a common source can be joined at a common target, vanishing of π_2 asserts that homotopies between common paths can be joined by a coherent higher homotopy.

The connection between higher homotopy groups and higher cohomology is made precise by the Hurewicz theorem

$$h : \pi_n(X_{\mathsf{adm}}, x_0) \to H_n(X_{\mathsf{adm}}; \mathbb{Z}),$$

which is an isomorphism when X_{adm} is $(n-1)$-connected. A complete admissibility theory would compute these invariants for each domain and compare them, establishing whether the admissibility structures in different domains are homotopy-equivalent or related only at the level of the fundamental group.

The Homotopy Type Theory program provides a natural foundational setting for this program. In HoTT, the identity type $\mathrm{Id}_A(x, y)$ is not a proposition but a type whose inhabitants are identity witnesses, that is, paths from x to y. The admissible identity path of the definition above is precisely such a term. Admissibility collapse corresponds to concatenation of identity witnesses, and the history H of a bubble in Spherepop is the accumulated data of all identity witnesses traversed. A future formalization of Spherepop may therefore find its natural home in cohesive homotopy type theory, which adds a geometric modality capable of expressing the spatial structure of bubble membranes.

F. Non-Vacuity Criterion and Boundaries of the Framework

This appendix develops the non-vacuity criterion for the admissibility framework in detail, addressing the objection that any constrained process might instantiate the generalized reduction operator.

F.1. The vacuity problem

The vacuity problem for any sufficiently general framework is that the framework's conditions may be satisfiable by every process in the intended domain, rendering the framework empirically empty. For the admissibility principle, the risk is real: Definition 11.1 specifies four conditions but these are individually weak enough that many processes satisfy them trivially. A thermostat satisfies all four under a liberal reading, yet it is clearly not the intended kind of system.

The resolution is to augment Definition 11.1 with conditions that are jointly strong enough to exclude trivial instantiations.

F.2. Strong admissibility: the augmented criterion

Definition F.1 (Strong admissibility)**.** A process $\mathcal{P}$ satisfies *strong admissibility* if it satisfies all four conditions of Definition 11.1 and the following three augmentations hold. First, the admissibility predicate Adm recursively restructures the future trajectory space: the set of admissible perturbations $\{\Delta : \mathsf{Adm}(\mathcal{B}_i, \Delta)\}$ changes in a history-dependent way not determined by $\mathcal{B}_i$ alone but by the full history $H(\mathcal{B}_i)$. Second, the invariant $\iota(\mathcal{B}_i)$ is transportable, meaning that downstream processes can use $\iota(\mathcal{B}_{i+1})$ as a constraint on their own admissibility without accessing the full history $H(\mathcal{B}_{i+1})$. Third, the system is path-sensitive: there exist pairs of histories H and H' leading to the same extensional state $\mathcal{B}$ but determining different future admissibility sets $\{\Delta : \mathsf{Adm}((\mathcal{B}, H), \Delta)\} \neq \{\Delta : \mathsf{Adm}((\mathcal{B}, H'), \Delta)\}$.

These three augmentations together constitute a significant strengthening. The first condition excludes all Markovian processes: in a Markov process the future transition kernel depends only on the current state, so the admissibility predicate cannot be recursively restructured by history. The second condition excludes processes that propagate invariants only internally without making them available for downstream use. The third condition excludes all path-independent processes, including thermostats, lookup tables, and arbitrary dissipative systems.

F.3. Canonical non-instances and their analysis

A thermostat fails the first condition because its future behavior is determined entirely by the current temperature reading, not the sequence of prior readings. The admissibility predicate does not change as a function of the history of prior adjustments. A thermostat instantiates weak admissibility at best, an output filter rather than a recursive state-space reorganizer.

A lookup table fails all three augmented conditions. Its future outputs are determined by current inputs alone, the table values are invariant under any sequence of queries, and the history of prior queries has no effect on future outputs.

A memoryless Markov chain fails the first and third conditions: future transition probabilities are determined by the current state alone, and no history-dependent restructuring of the state space occurs.

A physically dissipative system such as a glass shattering fails the second condition: no transportable invariant survives the transformation in a form usable by downstream processes. The shards do not carry the organizational invariant of the intact glass forward in any operationally meaningful way.

F.4. Biological development as a positive exemplar

The development of a multicellular organism from a fertilized egg is the canonical positive instance of strong admissibility; it is treated at length in Section 6 through the computational biology of multimodal lineage tracing (Wang, He, & Hu 2026). Here the essential formal features are summarized. At each stage, the admissibility predicate governing which cell fates are accessible is determined not merely by the cell's current gene expression state but by its full developmental history: the lineage of prior differentiation events, the morphogen gradients experienced, and the positional information accumulated through tissue interactions. This is recursive state-space restructuring in precisely the sense required by the first augmented condition.

The transportable invariant is positional and lineage identity: the information that a cell's descendants will use to determine their own admissible developmental options. This invariant is communicated through signaling molecules, transcription factor profiles, and epigenetic marks, making it externally accessible to neighboring cells without requiring access to the full developmental history. This satisfies the second augmented condition.

Finally, the system is strongly path-sensitive: two cells with identical current gene expression but different developmental histories will have different admissible futures even if their present transcriptional states are extensionally indistinguishable. This is precisely the third augmented condition. Biological development is therefore both a positive exemplar of strong admissibility and an existence proof that naturally occurring systems can satisfy all augmented conditions simultaneously.

G. The Three-Level Coarse-Graining: Syntax, Topology, Thermodynamics

G.1. The coarse-graining hierarchy

A persistent objection to the admissibility framework is that it appears to dissolve the distinctions between fundamentally different ontological categories. Symbolic rewriting systems, topological deformations, and thermodynamic evolutions are not members of the same mathematical universe, and identifying them risks category errors of a severe kind.

The resolution is that the admissibility framework does not identify these levels but asserts that the same admissibility grammar reappears at each level of coarse-graining. The progression is a hierarchy:

$$\text{syntax} \xrightarrow{\text{geometric realization}} \text{topology} \xrightarrow{\text{thermodynamic limit}} \text{thermodynamics},$$

in which each arrow represents a coarse-graining operation preserving admissibility structure while suppressing detail.

At the syntactic level, admissibility is determined by typing rules, capture-avoidance conditions, and the structure of explicit substitutions. The admissibility predicate is decidable: type-checking is computable, and the invariant is the typing judgement.

At the topological level, syntactic objects are replaced by geometric ones: terms become bubbles, substitution becomes membrane collapse, and the reduction relation becomes a continuous deformation of topological spaces. Admissibility is no longer decidable in general but it is geometrically natural. The invariant is topological: the homotopy type of the admissible locus, its fundamental group, the higher homotopy groups, and the sheaf cohomology.

At the thermodynamic level, topological objects are replaced by field configurations: bubbles become regions of the entropy-potential field, membrane collapse becomes gradient-constrained field relaxation, and reduction becomes a variational flow. Admissibility is a condition on the entropy production rate and the admissibility functional. The invariant is thermodynamic: the free energy, the entropy production history, and the admissible attractor basin.

G.2. The mesoscale necessity theorem

A structural consequence of this three-level hierarchy is the necessity of mesoscale organization. At the syntactic level, admissibility predicates operate on terms of finite size. At the thermodynamic level, admissibility conditions are expressed as integrals over extended regions. Neither level alone can support coherent admissibility transport across the full range of scales relevant to cognitive, physical, and semantic systems.

The topological level is the necessary intermediary. It provides geometric objects of finite extent but with non-trivial topological structure—bubble complexes, homotopy cells, sheaf stalks—that can simultaneously encode syntactic structure through their combinatorial topology and thermodynamic structure through their volume and boundary geometry.

The admissibility principle therefore predicts that stable computation, cognition, and physical organization require persistent mesoscale constraint structures. In turbulence, Reynolds stresses encode mesoscale admissibility information that is neither purely microscopic (individual fluid particle trajectories) nor purely macroscopic (mean flow statistics). In cognition, conceptual chunks are mesoscale attractors between neuronal microstates and explicit symbolic reasoning. In language models, embedding representations occupy the mesoscale between token-level statistics and global semantic coherence. Any theory that eliminates mesoscale organization should fail to produce coherent admissibility transport: this is a nontrivial and potentially falsifiable prediction of the framework.

G.3. The thermodynamic cost of syntactic irreversibility

The relationship between the syntactic and thermodynamic levels can be made partially precise through computational thermodynamics. Landauer's principle asserts that any logically irreversible computation must dissipate at least $k_B T \ln 2$ of energy per bit erased, where k_B is the Boltzmann constant and T is the temperature of the thermal environment.

In the admissibility framework, condition (iii) of Definition 11.1 asserts that each reduction strictly extends the history. This means that the admissibility framework is inherently logically irreversible in the Landauer sense: each collapse erases the pre-collapse state from the admissible futures, and this erasure has a minimum thermodynamic cost proportional to the information content of the erased pre-collapse state.

The total thermodynamic cost of a reduction sequence $\mathcal{B}_0 \to \mathcal{B}_1 \to \cdots \to \mathcal{B}_n$ is at least

$$W \geq n \cdot k_B T \ln 2 + \sum_{k=0}^{n-1} k_B T \cdot \log_2 |\mathcal{B}_k|,$$

where $|\mathcal{B}_k|$ is the information content of the pre-collapse state at step k. The RSVP entropy production limit $\dot{\Sigma}_{\text{crit}}$ can therefore be interpreted as an upper bound on the rate at which Landauer-irreversible operations can be performed: a system exceeding this limit produces entropy faster than its constraint geometry can absorb, leading to structural destabilization.

This connection identifies a specific and tractable research target for Layer V: deriving $\dot{\Sigma}_{\text{crit}}$ from Landauer's principle applied to admissibility-irreversible reductions. The derivation would require a precise information-theoretic account of the information content of admissibility regions and the Landauer cost of each collapse, which in turn requires the formal Spherepop specification identified in Section 3 as the primary open problem. The three formal gaps — Spherepop axiomatization, categorical embedding of $\mathcal{R}$, and thermodynamic derivation of $\dot{\Sigma}_{\text{crit}}$ — are therefore not independent: progress on any one of them directly advances the others.