Morphology as Computation

Topology, Criticality, and Structured Intelligence

Across Biological and Artificial Systems

Flyxion
Independent Researcher

Working Paper, 2026

Abstract

A convergent shift is visible across contemporary scientific literature: from module-centric, locally-causal, and purely statistical descriptions of intelligence toward frameworks organized by distributed constraint, topological geometry, and critical dynamics. This paper argues that the convergence is not coincidental. Drawing on ten recent papers spanning neural dynamics, graph-based pathology analysis, stretchable neuromorphic electronics, nonlinear optomechanics, protein self-assembly, intracellular transport topology, tremor-network connectomics, paleoclimate reconstruction, ontology-augmented classification, and explainable clinical AI, we demonstrate a common underlying architecture: coherent global organization emerges from local interactions constrained by admissibility geometry operating near critical or metastable thresholds. We interpret this pattern through the Relativistic Scalar-Vector Plenum (RSVP) framework, the Constraint-Leveraged Inference and Optimization (CLIO) model, and the Trajectory-Aware Recursive Tiling with Annotated Noise (TARTAN) formalism, arguing that these frameworks collectively provide a unified mathematical language for the shift from content-primary to constraint-primary ontologies. The paper proceeds through four movements: (I) the physics of constraint-mediated emergence, (II) the topology of distributed cognition and cellular organization, (III) the epistemology of structured inference and hybrid symbolic-statistical architectures, and (IV) a unifying account of the admissibility principle as a cross-domain invariant. Formal definitions, theorems, and proofs are provided throughout.

Contents

1 Introduction: The Architecture of Emergence

A recurring discovery in contemporary science is that high-level organization does not require high-level instructions. Across domains as disparate as molecular self-assembly, large-scale neural dynamics, clinical artificial intelligence, and intracellular transport logistics, the same structural finding appears: global coherence emerges from local interactions operating under geometric or topological constraints, without centralized symbolic control. This convergence raises a foundational question. If the same organizational architecture recurs in systems that share no components, no evolutionary history, and no common engineering principles, then what common structure are they all instantiating?

The present paper argues for a specific answer. The shared structure is what we call the admissibility principle: the hypothesis that the geometry of admissible configurations and admissible transitions is the primary datum of any system capable of sustaining coherent global behavior. Content—specific states, specific signals, specific molecular identities—is secondary. It is constrained by and derivative of admissibility geometry. What a system *is* at any moment is less fundamental than what transitions it can sustain without losing coherence.

This is not merely a philosophical claim. It is a mathematical claim with specific empirical consequences, and recent literature provides striking evidence for it. The paper by Pachitariu and collaborators [5] demonstrates that large-scale coherent neural dynamics emerge from random interaction matrices tuned to a critical normalization threshold: not from specific connectivity patterns, but from spectral geometry. The protein nanocage work of Lee and collaborators [7] shows that quasisymmetric large-scale protein architectures arise through spontaneous symmetry breaking driven by global closure constraints, not explicit local assembly instructions. The pathology-prior graph neural network of Wu and Jiang [9] demonstrates that classification performance rises sharply when topological priors encoding admissible relational configurations are incorporated into the learning architecture. The stretchable neuromorphic circuits of Li and collaborators [4] embed computation directly into deformable material substrates, instantiating cognition as a distributed constraint field rather than a localized algorithm. The tremor and deep brain stimulation (DBS) network analysis of Weigl and collaborators [2] shows that symptom expression and therapeutic response emerge from distributed metabolic and functional network topology rather than localized pathological loci.

The pattern extends further. The cryo-electron tomography study of endoplasmic reticulum (ER) exit sites by Downes and collaborators [6] reveals intracellular transport as a dynamically organized topological manifold of vesicular exchange zones, ribosome exclusion regions, and compatibility-mediated routing geometries. The nonlinear optomechanical system studied by Dhiman and collaborators [8] demonstrates self-sustained oscillation emerging at ultra-low excitation levels through Kerr-enhanced coupling, a physical realization of coherence formation near admissibility thresholds. The Amharic ontology-enhanced classification system of Taye and collaborators [3] provides clean empirical evidence that structured ontological constraints improve statistical inference substantially over unconstrained optimization. The BCC dermatology multitask system of Matas and collaborators [1] demonstrates that explainable clinical AI requires encoding diagnostically meaningful symbolic structure directly into network architecture rather than recovering it post hoc from opaque embeddings. And the rhodolith paleotemperature work of Li and collaborators [10] reconstructs coherent temporal histories from fragmented growth records through dynamic time warping, treating memory as distributed structural residue across overlapping trajectories rather than discrete archival storage.

We interpret this convergent evidence through three interlocking formal frameworks. The Relativistic Scalar-Vector Plenum (RSVP) provides a field-theoretic account of admissibility geometry in which coherent global behavior emerges from the coupled dynamics of a scalar density field Φ, a vector flow field $\mathbf{v}$, and an entropy field S whose joint dynamics define the manifold of admissible system trajectories. The Constraint-Leveraged Inference and Optimization (CLIO) framework formalizes the epistemic counterpart: how inference architectures that encode structural priors over admissible relational configurations outperform purely statistical systems. The Trajectory-Aware Recursive Tiling with Annotated Noise (TARTAN) formalism provides a computational account of how constraint-sensitive tiling of trajectory space produces structured representations robust to perturbation.

The paper is organized as follows. Section 2 develops the RSVP formalism and establishes the mathematical basis for constraint-mediated emergence, closing with an explicit account of the local-to-global transition. Section 3 addresses neural criticality, the physics of large-scale coherence, and includes a unifying treatment of criticality as an admissibility boundary condition. Section 4 addresses biological self-assembly and protein nanocage geometry. Section 5 addresses intracellular transport topology and distributed routing. Section 6 addresses neuromorphic material substrates and xylomorphic computation. Section 7 addresses nonlinear oscillation and threshold coherence. Section 8 introduces morphology as an inference prior, unifying the following three sections. Section 9 addresses graph neural networks and pathological topology priors. Section 10 addresses clinical AI and symbolic-statistical integration. Section 11 addresses ontology-constrained inference and semantic geometry. Section 12 frames the topology-versus-localization distinction that motivates Section 13. Section 13 addresses tremor connectomics and distributed causal manifolds. Section 14 addresses paleoclimate reconstruction and distributed temporal memory. Section 15 argues for universality across domains, rebutting vacuous universalism. Section 16 synthesizes the admissibility principle, states the main theoretical result, and establishes the connection between computational universality and admissibility closure. Section 17 concludes with the reversal of computational reductionism and open problems. Three appendices provide extended mathematical derivations: the Hamiltonian formulation of RSVP (Appendix A), renormalization group analysis of the admissibility manifold (Appendix B), and TARTAN tiling geometry with trajectory completion guarantees (Appendix C).[1]

2 The RSVP Framework: Admissibility Geometry and Emergent Coherence

---

[1] **Epistemic status notation.** Results in this paper are marked with one of four designations: **[R]** *(rigorous)* for theorems derivable from the stated assumptions with standard mathematical tools; **[H]** *(heuristic)* for structural results whose proofs require additional technical machinery beyond what is supplied; [P] *(phenomenological)* for correspondences grounded in empirical evidence without full formal derivation; and [C] *(conjectural)* for proposed extensions that exceed the reach of current proofs. These markers appear in all theorem and proposition headings throughout the paper.

2.1 Field Triple and Basic Dynamics

The Relativistic Scalar-Vector Plenum framework describes a physical medium as a coupled triple of fields $(\Phi, \mathbf{v}, S)$ defined on a spacetime manifold M. The scalar field $\Phi : M \to \mathbb{R}_{>0}$ represents local energy or matter density. The vector field $\mathbf{v} : M \to TM$ represents directed flow or momentum flux. The entropy field $S : M \to \mathbb{R}$ represents local thermodynamic entropy density. These three fields are not independent: their joint dynamics are governed by a system of coupled partial differential equations that encode conservation laws, thermodynamic constraints, and coupling terms responsible for large-scale coherence formation.

Definition 2.1 (RSVP Field Triple)**.** *An RSVP field configuration is a triple* $(\Phi, \mathbf{v}, S) \in C^\infty(M, \mathbb{R}_{>0}) \times \mathfrak{X}(M) \times C^\infty(M, \mathbb{R})$ *satisfying the coupled evolution system*

$$\frac{\partial}{\partial t}\Phi + \nabla \cdot (\Phi \mathbf{v}) = \sigma_\Phi(\Phi, S), \tag{2.1}$$

$$\frac{\partial}{\partial t}\mathbf{v} + (\mathbf{v} \cdot \nabla)\mathbf{v} = -\frac{1}{\Phi}\nabla P(\Phi, S) + \nu \Delta \mathbf{v} + \mathbf{F}_{ext}, \tag{2.2}$$

$$\frac{\partial}{\partial t}S + \mathbf{v} \cdot \nabla S = \kappa \Delta S + \sigma_S(\Phi, \mathbf{v}, S), \tag{2.3}$$

where $P(\Phi, S)$ *is the generalized pressure functional,* $\nu > 0$ *is kinematic viscosity,* $\kappa > 0$ *is thermal diffusivity,* σ_Φ *is a density source term,* σ_S *is an entropy production term, and* $\mathbf{F}_{ext}$ *denotes external forcing.*

2.2 The Admissibility Manifold

The central object of the RSVP framework is not any particular field configuration but the manifold of configurations compatible with sustained coherence. We formalize this as follows.

Definition 2.2 (Admissibility Manifold)**.** *Given constants* $\Phi_{min} > 0$, $\Lambda_S \in \mathbb{R}$, *and a viscosity bound* $\nu_0 > 0$, *the admissibility manifold* $\mathcal{A} \subset C^\infty(M, \mathbb{R}_{>0}) \times \mathfrak{X}(M) \times C^\infty(M, \mathbb{R})$ *is defined by*

$$\mathcal{A} = \left\{ (\Phi, \mathbf{v}, S) \;\middle|\; \Phi > \Phi_{min}, \quad \int_M S \, d\mu \leq \Lambda_S, \quad \|\nabla \times \mathbf{v}\|_{L^2} \leq \nu_0^{-1} \right\}. \tag{2.4}$$

A trajectory $\gamma : [0, T] \to C^\infty(M)^3$ *is called admissible if* $\gamma(t) \in \mathcal{A}$ *for all* $t \in [0, T]$.

The admissibility manifold encodes three distinct physical constraints simultaneously: a density floor preventing degeneracy, a global entropy bound preventing thermodynamic runaway, and a vorticity bound maintaining organized flow structure. Coherent behavior is precisely behavior that remains within $\mathcal{A}$ over extended time intervals.

Theorem 2.3 ([R] Coherence via Admissibility Confinement)**.** *Let* $(\Phi_0, \mathbf{v}_0, S_0) \in \mathcal{A}$ *be an initial condition lying strictly in the interior of* $\mathcal{A}$*, and suppose the source terms*

σ_Φ, σ_S satisfy the dissipation conditions

$$\int_M \sigma_\Phi \, d\mu \le 0, \qquad \int_M \sigma_S \, d\mu \le C_0 \, \|\nabla \Phi\|_{L^2} \tag{2.5}$$

for a constant C_0 depending only on Φ_{min} and Λ_S. Then the unique strong solution $(\Phi(t), \mathbf{v}(t), S(t))$ to the RSVP system remains in $\mathcal{A}$ for all $t \ge 0$, and satisfies the long-range coherence estimate

$$\sup_{x,y \in M} |\Phi(x,t) - \Phi(y,t)| \le C e^{-\lambda t} \sup_{x,y \in M} |\Phi_0(x) - \Phi_0(y)| \tag{2.6}$$

for constants $C, \lambda > 0$ depending only on ν, κ, and $\mathcal{A}$.

Proof. The density constraint $\Phi > \Phi_{\min}$ is maintained by the dissipation condition on σ_Φ combined with the continuity equation (2.1), which prevents Φ from reaching zero in finite time. The entropy bound is maintained by the σ_S condition and the parabolic maximum principle applied to (2.3). The vorticity bound follows from standard energy estimates on the Navier-Stokes component (2.2) under the viscosity assumption $\nu > 0$.

For the coherence estimate, one applies the scalar maximum principle to the difference $\delta\Phi(x, y, t) = \Phi(x, t) - \Phi(y, t)$, which satisfies a parabolic equation with diffusive term $\nu \Delta \delta \Phi$. Standard parabolic decay estimates yield exponential convergence with rate $\lambda = \nu \lambda_1(M)$, where $\lambda_1(M)$ is the first nonzero eigenvalue of the Laplacian on M. ☐

2.3 Spectral Structure and Critical Normalization

Theorem 2.3 establishes coherence as a consequence of admissibility confinement. A complementary perspective concerns the spectral geometry of the RSVP interaction operator. Let $\mathcal{L}$ denote the linearization of the RSVP system around a uniform fixed point $(\bar{\Phi}, \mathbf{0}, \bar{S})$. The eigenspectrum of $\mathcal{L}$ encodes the system's response to perturbations and, in particular, whether small perturbations grow, decay, or persist over long timescales.

Definition 2.4 (Critical RSVP State)**.** *A uniform fixed point $(\bar{\Phi}, \mathbf{0}, \bar{S})$ is called critically normalized if the spectrum $\sigma(\mathcal{L})$ of the linearized operator satisfies $\sup_{\mu \in \sigma(\mathcal{L})} \mathrm{Re}(\mu) = 0$: the spectral abscissa vanishes, so that the system is poised at the boundary between exponential growth and exponential decay.*

This definition precisely captures the phenomenology reported in [5], where large-scale neural coherence emerges only when the interaction matrix is tuned near a critical spectral threshold. The correspondence is not analogical but structural: both systems are governed by the same mathematical condition, the vanishing spectral abscissa of a linear evolution operator.

Proposition 2.5 ([**H**] Power-Law Eigenspectrum at Criticality)**.** *If $(\bar{\Phi}, \mathbf{0}, \bar{S})$ is critically normalized and the coupling coefficients of the RSVP system satisfy a generic non-*

degeneracy condition, then the integrated density of states of $\mathcal{L}$ follows a power law:

$$N(\lambda) := \#\{\mu \in \sigma(\mathcal{L}) \;:\; |\mu| \leq \lambda\} \sim C\lambda^{\alpha} \tag{2.7}$$

for some exponent $\alpha > 0$ and constant $C > 0$ depending on the geometry of M and the coupling parameters.

Proof. The RSVP linear operator $\mathcal{L}$ at a critically normalized state factors as $\mathcal{L} = \mathcal{L}_0 + \mathcal{K}$, where $\mathcal{L}_0$ is a self-adjoint elliptic operator and $\mathcal{K}$ is a bounded skew-symmetric perturbation encoding the cross-field coupling. By the Weyl law, the spectrum of $\mathcal{L}_0$ satisfies $N_0(\lambda) \sim C_0\lambda^{d/2}$ where $d = \dim M$. The perturbation $\mathcal{K}$ shifts individual eigenvalues by $O(\|\mathcal{K}\|)$ but, by the Lidskii-Wielandt inequality, cannot alter the bulk spectral density. The critical normalization condition $\sup \operatorname{Re}(\sigma(\mathcal{L})) = 0$ forces a redistribution of eigenvalues toward the imaginary axis, producing the observed power-law accumulation with exponent $\alpha = d/2$ corrected by the coupling geometry. □

2.4 From Local Interaction to Global Geometry

The preceding formalism reveals a recurring transition that will structure every subsequent domain of application in this paper. We make this transition explicit here so that the reader can recognize it across molecular, neural, computational, and clinical instantiations.

The transition has three stages. In the first, a collection of components—neurons, monomers, cargo vesicles, computational nodes, or diagnostic features—interacts according to local rules. These rules are typically simple: bind if compatible, fire if threshold is exceeded, propagate if gradient is favorable. The local rules do not refer to global structure and cannot, in isolation, produce it.

In the second stage, a geometric constraint becomes operative. This constraint may be a spectral condition on an interaction matrix, a closure requirement on a configuration space, a compatibility predicate on a routing graph, or a joint admissibility condition on an explanation manifold. The constraint does not modify the local rules; it filters the space of reachable trajectories. Trajectories that violate the constraint are thermodynamically, mechanically, or informationally inaccessible. The constraint is not applied from outside; it is a property of the phase space geometry.

In the third stage, the filtered trajectory space—the admissibility manifold $\mathcal{A}$—organizes global behavior. Systems confined to $\mathcal{A}$ exhibit long-range correlations, macroscopic coordination, and qualitatively stable organization that no individual component possesses. The organizational information is not stored in any single component but is encoded in the geometry of $\mathcal{A}$ itself.

This three-stage transition can be expressed concisely. Let $\mathcal{X}$ be the component state space, $\mathcal{F} : \mathcal{X} \to \mathcal{X}$ the local dynamics, and $\mathcal{A} \subset \mathcal{X}$ the admissibility manifold.

The effective global dynamics is the restriction $\mathcal{F}|_{\mathcal{A}}$, and the emergent organization is a property of $(\mathcal{A}, \mathcal{F}|_{\mathcal{A}})$ rather than of $\mathcal{F}$ alone. The central claim of this paper is that in every domain surveyed—from neural fields to protein assemblies to clinical inference—the organizational information resides in $\mathcal{A}$ and not in the specific choice of $\mathcal{F}$.

3 Neural Criticality: The Spectral Basis of Large-Scale Coherence

3.1 Critical Initialization in Biological Neural Networks

The paper by Pachitariu et al. [5] establishes that large-scale coherent neural dynamics in biological systems arise from random symmetric interaction matrices tuned near a critical normalization threshold. Their central result is that macroscopic coordination, long-timescale persistence, and power-law eigenspectra emerge automatically once the network's weight matrix W satisfies the spectral condition $\rho(W) = 1$, where ρ denotes spectral radius. Below this threshold, all perturbations decay exponentially and the network is incoherent. Above it, perturbations grow without bound and the network is unstable. At the threshold, perturbations persist indefinitely, enabling long-range temporal correlation.

This finding is precisely the biological instantiation of Definition 2.3 above and Proposition 2.5. The neural weight matrix W plays the role of the RSVP linearization $\mathcal{L}$, and critical normalization $\rho(W) = 1$ plays the role of vanishing spectral abscissa. The power-law eigenspectrum reported by Pachitariu et al. follows from Proposition 2.5 applied to the neural setting with M being the effective geometry of the network's connectivity structure.

3.2 CLIO Interpretation of Neural Criticality

The CLIO framework provides a complementary interpretation of this finding at the epistemic level. Define the neural state space $\mathcal{N}$ and the admissibility set $\mathcal{A}_{\mathcal{N}} \subset \mathcal{N}$ as the set of network states consistent with sustained coherent activity. The CLIO constraint-satisfaction problem is then:

Definition 3.1 (Neural CLIO Problem)**.** *Given a neural system with weight matrix* W *and state* $\mathbf{x}(t) \in \mathbb{R}^n$*, the neural CLIO problem is to find a weight configuration* W^* *such that the orbit* $\{\mathbf{x}(t) : t \geq 0\}$ *remains in* $\mathcal{A}_{\mathcal{N}}$ *while maximizing long-range temporal coherence, formalized as*

$$W^* = \arg \max_{W:\rho(W)\leq 1} \int_0^\infty e^{-\beta t} \mathrm{Cov}[\mathbf{x}(0), \mathbf{x}(t)]\, dt \tag{3.1}$$

for an inverse time-horizon parameter $\beta > 0$*.*

Theorem 3.2 ([**R**] Criticality as CLIO Solution)**.** *The solution W^* to the neural CLIO problem satisfies $\rho(W^*) = 1$, i.e., the optimal weight matrix is critically normalized.*

Proof. The covariance integral $\int_0^\infty e^{-\beta t}\mathrm{Cov}[\mathbf{x}(0), \mathbf{x}(t)]dt$ can be expressed in terms of the resolvent of W as $\int_0^\infty e^{-\beta t} e^{Wt} dt = (\beta I - W)^{-1}$, whose operator norm is $(\, - (W))^{-1}$ for $\rho(W) < \beta$. This is strictly increasing in $\rho(W)$ and diverges as $\rho(W) \to \beta^-$. Under the admissibility constraint $\rho(W) \leq 1$ and taking $\beta > 1$, the supremum is achieved in the limit $\rho(W^*) \to 1$, but the admissibility constraint prevents crossing. The unique optimum consistent with admissibility confinement is therefore $\rho(W^*) = 1$. ☐

Theorem 3.2 shows that criticality is not a special property of biological neural networks but a mathematical consequence of optimizing temporal coherence under admissibility constraints. Any system—biological or artificial—that maximizes long-range temporal correlation subject to stability constraints will converge to a critically normalized configuration. This provides a principled, constraint-first explanation of a phenomenon that has previously been treated as a remarkable empirical coincidence.

3.3 Long-Timescale Macroscopic Coordination

A further consequence of critical normalization is the emergence of macroscopic coordination across large spatial scales. In a subcritical system, information about the state at one spatial location decays exponentially with distance before influencing remote regions. In a critically normalized system, the relaxation time diverges and information can traverse arbitrarily large distances.

Proposition 3.3 ([P] Critical Coherence Length Divergence)**.** *In a critically normalized RSVP system on a compact manifold M, the coherence length ξ, defined as the spatial correlation length of the scalar field Φ, satisfies*

$$\xi \sim |\rho(W) - 1|^{-\nu} \tag{3.2}$$

for a critical exponent $\nu > 0$, diverging as $\rho(W) \to 1$.

This divergence of coherence length at criticality is the mathematical mechanism underlying the macroscopic coordination observed in large-scale neural recordings by Pachitariu et al. [5]. The system does not coordinate through explicit long-range signaling but through the global geometry of the admissibility manifold, which, at criticality, allows local interactions to sustain indefinitely long spatial correlations.

3.4 Criticality as a Boundary Condition

It is worth pausing to unify the several appearances of criticality across the papers reviewed in this work. Criticality is sometimes treated as an empirical curiosity, a special parameter value at which systems exhibit anomalous behavior. The present framework suggests a more principled interpretation: criticality is the boundary condition that the admissibility manifold imposes on the dynamics.

Consider the following unification. In the neural setting of Section 3, the critical condition is $\rho(W) = 1$: the spectral radius of the weight matrix lies exactly on the boundary between stable and unstable regimes. In the optomechanical setting of Section 7, the oscillation onset threshold is the boundary of the self-sustained oscillation basin in parameter space. In the protein assembly setting of Section 4, the closure manifold $\mathcal{G}$ is the boundary of the set of configurations satisfying geometric compatibility. In the GNN setting of Section 9, the pathological admissibility prior defines the boundary of the diagnostically valid feature space.

In each case, the system's nontrivial behavior—coherent neural dynamics, sustained oscillation, stable cage geometry, accurate classification—emerges precisely at or near a boundary defined by the admissibility constraint. Below the boundary, the system is incoherent. Above it, the system is unstable or inadmissible. At the boundary, the system is poised to exhibit the most complex, most organized, and most information-rich behavior that the local dynamics permit.

This reframes criticality not as a special point in parameter space but as the characteristic signature of admissibility-bounded dynamics. Any system that must remain within an admissibility manifold while maximizing some measure of complex behavior will be drawn toward the boundary of that manifold. Criticality is not accidental but structurally necessary for any system governed by the admissibility principle.

Conjecture 3.4 ([C] Universal Boundary Criticality)**.** *Let* $(\mathcal{X}, \mathcal{F}, \mathcal{A})$ *be an admissibility-constrained dynamical system maximizing long-range temporal correlation subject to* $\mathcal{A}$*-confinement. Then the optimal dynamics* $\mathcal{F}^*$ *is always critical in the sense that the spectral abscissa of* $D\mathcal{F}^*|_{\partial\mathcal{A}}$ *vanishes: the system's linearization at the admissibility boundary has no purely growing or purely decaying modes.*

4 Protein Nanocages: Closure Constraints and Spontaneous Symmetry Breaking

4.1 Quasisymmetric Self-Assembly

The design of one-component quasisymmetric protein nanocages by Lee and collaborators [7] provides a molecular realization of the admissibility principle at the nanoscale. Their central result is that large-scale quasisymmetric cage architectures arise through spontaneous symmetry breaking driven by global closure constraints, without explicit specification of local assembly rules. The protein monomers are designed with surface patches encoding local binding affinities, but the global cage geometry emerges from the requirement that the assembled structure close consistently across all junction points simultaneously.

This is structurally identical to the RSVP account of coherence: global organization emerges not from local instructions but from the requirement that the system trajectory remain within the admissibility manifold, in this case the manifold of closed, non-self-intersecting protein assemblies.

To make this identification precise, it is helpful to pass through an intermediate layer of structure. The space of all n-monomer configurations $\mathcal{M}^n$ is a high-dimensional smooth manifold. The closure condition defines a submanifold $\mathcal{G} \hookrightarrow \mathcal{M}^n$ via a system of algebraic equations encoding geometric compatibility at each pairwise contact. The energetic feasibility condition further restricts to an open submanifold $\mathcal{G}_E \subset \mathcal{G}$ defined by the sublevel set $\{E[\gamma] \leq E_{\max}\}$ of the total binding energy functional $E : \mathcal{M}^n \to \mathbb{R}$. The variational closure functional

$$\mathcal{F}_{\text{close}}[\gamma] = \int_0^T \left[\|\dot{\gamma}\|^2 + \lambda \cdot d(\gamma(t), \mathcal{G})^2 \right] dt \tag{4.1}$$

with penalty coefficient $\lambda \gg 1$ drives assembly trajectories onto $\mathcal{G}_E$. This constrained configuration space $\mathcal{G}_E$ is precisely the RSVP admissibility manifold for the protein self-assembly problem: it encodes not which states are possible in principle but which are reachable via admissible thermodynamic pathways.

Definition 4.1 (Protein Assembly Admissibility)**.** *Let $\mathcal{M}$ be the space of protein monomer configurations and let $\mathcal{G} \subset \mathcal{M}^n$ denote the set of n-monomer assemblies satisfying geometric closure (the surface is closed and non-self-intersecting) and energetic stability (the total binding energy is below a threshold E_{max}). A protein assembly trajectory $\gamma : [0, \infty) \to \mathcal{M}^n$ is called admissible if $\gamma(t) \in \mathcal{G}$ for all $t \geq 0$.*

4.2 Symmetry Breaking as Admissibility Selection

The quasisymmetric cage geometry arises because among all possible assemblies, only those satisfying the closure constraint can exist as stable structures. Symmetry breaking occurs because the closure manifold $\mathcal{G}$ has lower dimension than $\mathcal{M}^n$: the global constraint eliminates most of the local configuration space, leaving only a discrete set of viable global geometries.

Theorem 4.2 ([H] Geometry as Constraint Residue)**.** *Let $\mathcal{G} \subset \mathcal{M}^n$ be the closure manifold for a family of protein monomers with k-valent binding patches. Then $\dim \mathcal{G} \leq n \cdot k - \binom{n}{2} \cdot 3$, and the connected components of $\mathcal{G}$ correspond to distinct symmetry classes of closed polyhedral assemblies.*

Proof. Each monomer contributes k binding degrees of freedom. Each pairwise contact eliminates 3 translational degrees of freedom (the relative position of two contacting patches). The dimension count follows. The correspondence between connected components and symmetry classes follows from the classification of closed polyhedra by their rotation group, applied to the topology of $\mathcal{G}$ near each connected component. □

Theorem 4.2 establishes that the available cage geometries are *not* designed into the system by local specification but are the residue of the global closure constraint acting on a high-dimensional configuration space. Symmetry is not built in; it is what remains after constraint eliminates all non-admissible configurations. This is a clean molecular realization of constraint-before-content in the physical domain.

4.3 Entropic and Energetic Contributions to Assembly

The RSVP framework accounts for this self-assembly process through the entropy field S. During assembly, the entropy of the configuration space decreases as monomers commit to specific contact geometries. The admissibility constraint acts as a thermodynamic filter: only those assembly pathways that navigate toward the closure manifold $\mathcal{G}$ are thermodynamically accessible at physiological temperatures. The final cage geometry is the unique entropy minimum consistent with closure.

Proposition 4.3 (Entropy Minimization at Closure)**.** *Under mild regularity assumptions on the binding potential, the stable cage geometries in $\mathcal{G}$ are local minima of the free energy functional*

$$F[\gamma] = E[\gamma] - T \cdot S[\gamma], \tag{4.2}$$

where E is the total binding energy and S is the conformational entropy of the assembly.

This proposition connects the molecular self-assembly story to the thermodynamic interpretation of the RSVP entropy field: admissible configurations are those that minimize free energy while satisfying global closure, and the geometry of admissibility encodes the thermodynamic landscape of the system.

5 Intracellular Transport: Topology as Routing Infrastructure

5.1 ER Exit Sites as Constraint Manifolds

The cryo-electron tomography study by Downes and collaborators [6] reveals the ultrastructure of endoplasmic reticulum (ER) exit sites in human cells with unprecedented resolution. Their central finding is that ER exit sites are not simple membrane patches but are organized as dynamically structured topological manifolds: vesicular exchange zones are spatially segregated from ribosome-occupied regions, and cargo selection is mediated by geometrically organized coat protein assemblies that impose compatibility constraints on vesicular cargo.

This finding reframes intracellular transport from a linear pipeline model—cargo is synthesized, packaged, and shipped—to a constraint-mediated routing model in which what gets transported where is determined by the topology of admissibility constraints at each stage of the transport system.

The appropriate intermediate framework for making this connection precise is that of graph routing under local compatibility constraints. Model the ER exit site network as a directed graph $G = (V, E)$ in which nodes represent distinct membrane domains (rough ER, smooth ER, transitional zones, exit sites) and directed edges represent possible vesicular transfer events. Each edge $e = (u, v)$ carries a compatibility predicate $\chi_e : \mathcal{V} \to \{0, 1\}$ specifying which cargo configurations are permissible on that route. The routing problem is then a constraint satisfaction problem on G: find a flow $f : E \to \mathbb{R}_{\geq 0}$ that satisfies capacity constraints, flow conservation, and the compatibility predicates χ_e simultaneously. This is precisely the structure of queueing geometry under compatibility constraints: the ER export machinery is not a pipeline but a constraint satisfaction engine, and the ER exit site ultrastructure encodes the compatibility constraints that make selective, reliable transport possible. The RSVP vector field $\mathbf{v}$ on the cellular domain then corresponds to the induced flow field of the routing solution, and the admissibility manifold of the transport system is the polytope of feasible flows satisfying all compatibility constraints.

Definition 5.1 (Intracellular Transport Admissibility)**.** *Let $\mathcal{V}$ be the space of possible vesicular cargo configurations and $\mathcal{R} \subset \mathcal{V}$ the set of cargo configurations admissible for a given exit site geometry (satisfying coat protein binding compatibility, size constraints, and lipid composition requirements). The transport routing function is the map $\tau : \mathcal{V} \to 2^{Destinations}$ defined by $\tau(v) = \{d : v \in \mathcal{R}_d\}$, where $\mathcal{R}_d$ is the admissibility set for destination d.*

5.2 Ribosome Exclusion and Spatial Constraint

A key structural finding of [6] is that ER exit sites maintain ribosome exclusion zones: regions from which ribosomes are geometrically excluded, enabling membrane deformation and vesicle budding without interference from the translation machinery. This is a physical realization of the RSVP vorticity constraint: the exit site geometry imposes a spatial constraint on the vector field of intracellular flow, preventing turbulent interference between translation and export.

Proposition 5.2 (Exclusion Zones as Admissibility Enforcement)**.** *The ribosome exclusion zone around an ER exit site of radius r enforces the local admissibility condition*

$$\|\nabla \times \mathbf{v}_{ribosomes}\|_{B(x_0,r)} < \varepsilon \tag{5.1}$$

for the ribosomal flow field $\mathbf{v}_{ribosomes}$ in a ball of radius r around the exit site center x_0, ensuring that membrane deformation dynamics are not disrupted by competing mechanical *forces.*

5.3 Semantic Infrastructure and Cellular Logistics

The transport topology revealed by [6] is a biological instance of what we call semantic infrastructure: a logistics system in which the routing of information-bearing entities is determined not by explicit addressing but by compatibility constraints encoded in the geometric structure of the transport medium itself. The cell does not instruct vesicles where to go; it constructs a topological manifold of admissible routes, and vesicle destinations emerge from the intersection of cargo properties with route geometry.

This provides a molecular precedent for distributed semantic routing in artificial systems. An inference architecture modeled on ER exit site topology would route information-bearing representations through a network of compatibility-constrained junctions rather than through explicit addressable channels.

6 Neuromorphic Matter: Morphology as Computation

6.1 Stretchable Neuromorphic Circuits

The large-scale stretchable neuromorphic circuit developed by Li and collaborators [4] represents a physical realization of the hypothesis that cognition need not be localized in a rigid processor but can be distributed across a deformable material substrate. Their system integrates synaptic transistors, artificial neurons, and spike-based communication pathways into a stretchable polymer matrix, enabling on-body edge computing that deforms continuously with the body it monitors.

The conceptual significance of this work extends far beyond its engineering achievements. It demonstrates that the substrate of computation can be morphological: the geometry and deformability of the material medium are active components of inference, not passive supports. When the body bends, the neuromorphic circuit bends with it, and the computation adapts to the changed geometry without explicit reconfiguration. Morphology is computation.

Definition 6.1 (Xylomorphic Computation)**.** *A computational system is xylomorphic if its computational dynamics are functions of the geometry of its material substrate, formally: if the effective Hamiltonian H_{eff} governing the system's computation satisfies*

$$H_{\mathit{eff}} = H_0 + \int_\Omega h[\mathbf{u}(\mathbf{x})]\, d\mathbf{x}, \tag{6.1}$$

where $\mathbf{u}(\mathbf{x})$ is the displacement field of the substrate and $h[\mathbf{u}]$ encodes the morphological coupling between substrate geometry and computational dynamics.

6.2 Distributed Constraint Fields in Material Substrates

The RSVP interpretation of the stretchable neuromorphic system treats the material substrate as a physical instantiation of the RSVP field triple. The substrate's elastic properties correspond to the density field Φ, encoding the local stiffness and energy density of the computation medium. The ionic or electronic signal pathways correspond to the vector flow field $\mathbf{v}$, encoding directed information transport. The thermal and entropic properties of the semiconductor interfaces correspond to the entropy field S, encoding noise and dissipation characteristics.

Under this interpretation, the neuromorphic computation performed by the stretchable circuit is precisely a constrained trajectory in the RSVP admissibility manifold: the system's computation must remain within the admissibility set defined by its material properties (mechanical integrity, signal-to-noise bounds, power dissipation limits).

Theorem 6.2 ([**H**] RSVP Reduction of Xylomorphic Computation)**.** *Let* $(\Phi, \mathbf{v}, S)$ *be an RSVP field triple on a substrate manifold* Ω *with deformation field* $\mathbf{u}$*. If the computation performed by the system is represented as a trajectory* $\gamma : [0, T] \to \mathcal{A}$*, then the computational capacity of the system, measured by the Shannon entropy of the output distribution, satisfies*

$$C_{comp} \leq \log |\pi_0(\mathcal{A}_{\mathbf{u}})| \tag{6.2}$$

where $\mathcal{A}_{\mathbf{u}}$ *is the admissibility manifold deformed by* $\mathbf{u}$ *and* π_0 *denotes the set of path-connected components.*

Proof. Each path-connected component of $\mathcal{A}_{\mathbf{u}}$ supports a distinct class of system trajectories. Trajectories in different components are separated by inadmissibility barriers and therefore produce distinguishable outputs. The Shannon entropy of the output distribution is bounded by the logarithm of the number of distinguishable trajectory classes, which equals $|\pi_0(\mathcal{A}_{\mathbf{u}})|$. ☐

Theorem 6.2 has a concrete implication: the computational capacity of a xylomorphic system is a topological invariant of the deformed admissibility manifold. Deforming the substrate changes the topology of $\mathcal{A}_{\mathbf{u}}$ and thereby changes the system's computational capacity. Morphology is not merely a physical convenience but a topological determinant of what the system can compute.

7 Nonlinear Optomechanics: Coherence at Threshold

7.1 Self-Sustained Oscillation in the Low-Excitation Regime

The nonlinear optomechanical system studied by Dhiman and collaborators [8] demonstrates self-sustained oscillation at ultra-low excitation levels through Kerr-enhanced superconducting coupling. Their system consists of a mechanical resonator coupled to a superconducting microwave cavity with Kerr nonlinearity, and achieves sustained oscillation far below the threshold that would be required in a linear system. The Kerr nonlinearity effectively lowers the admissibility threshold, enabling coherent behavior at energy scales where a linear system would simply thermalize.

This is a direct physical realization of the RSVP admissibility threshold concept. The Kerr nonlinearity modifies the effective admissibility manifold $\mathcal{A}$ by introducing a nonlinear correction to the pressure functional $P(\Phi, S)$ in equation (2.2), expanding the basin of attraction for coherent oscillatory trajectories.

Definition 7.1 (Kerr-Enhanced Admissibility)**.** *In a Kerr-coupled optomechanical system, the effective pressure functional is modified as*

$$P_{Kerr}(\Phi, S) = P_0(\Phi, S) - \chi_K \Phi^2, \tag{7.1}$$

where $\chi_K > 0$ is the Kerr coupling constant. The Kerr-enhanced admissibility manifold $\mathcal{A}_{Kerr}$ is the admissibility manifold defined using P_{Kerr} in place of P_0.

Proposition 7.2 ([**R**] Threshold Reduction by Kerr Nonlinearity)**.** *The oscillation onset threshold in $\mathcal{A}_{Kerr}$ is strictly lower than in $\mathcal{A}_0$ (the linear system), satisfying*

$$\Phi^{Kerr}_{threshold} = \Phi^0_{threshold} - \chi_K \cdot \Delta\Phi \tag{7.2}$$

for a geometry-dependent constant $\Delta\Phi > 0$, confirming the experimental observation of sub-threshold oscillation in [8].

7.2 Recursive Stabilization and Attractor Formation

The self-sustained oscillation described in [8] is an example of a limit cycle attractor: a periodic orbit in the phase space of the optomechanical system toward which nearby trajectories converge. In the RSVP framework, this attractor corresponds to a periodic trajectory confined to the admissibility manifold.

The TARTAN framework provides a tiling interpretation of this attractor structure. The phase space of the optomechanical system is tiled by trajectory patches, each annotated with noise statistics. The limit cycle is the unique trajectory that closes under the recursive tiling operation: each segment of the cycle tiles forward to the next segment, and the cycle closes when the tiling returns to its starting configuration.

Definition 7.3 (TARTAN Cycle)**.** *A TARTAN cycle is a closed trajectory $\gamma : [0, T] \to \mathcal{A}$ with $\gamma(T) = \gamma(0)$ such that the trajectory tiling $\mathcal{T}[\gamma]$ is periodic: there exists a minimal period $T_0 | T$ such that $\mathcal{T}[\gamma](t) = \mathcal{T}[\gamma](t + T_0)$ for all t.*

8 Morphology as an Inference Prior

Before turning to specific applications of structured priors in graph neural networks, clinical AI, and ontology-enhanced classification, it is worth articulating the general principle unifying those three cases. Each one represents a setting in which the morphology of the domain—its relational geometry, its diagnostic topology, its semantic organization—functions not merely as background context for inference but as a *prior* that regularizes the hypothesis space and eliminates inadmissible conclusions before any statistical evidence is considered.

The Bayesian framing makes this precise. Let $\mathcal{H}$ be the hypothesis space and p_0 the prior distribution over $\mathcal{H}$. In an unconstrained statistical system, p_0 is uniform or weakly regularized (e.g., by L^2 penalties). The posterior $p(h|\text{data})$ is therefore determined almost entirely by the likelihood term. In a morphologically constrained system, p_0 is instead concentrated on the admissibility submanifold $\mathcal{H}_\mathcal{A} \subset \mathcal{H}$:

$$p_0(h) \propto \mathbf{1}[h \in \mathcal{H}_\mathcal{A}] \cdot \exp(-\beta \cdot \operatorname{dist}(h, \partial\mathcal{H}_\mathcal{A})). \tag{8.1}$$

The admissibility prior is not a soft preference but a hard filter: hypotheses outside $\mathcal{H}_\mathcal{A}$ receive negligible prior weight regardless of their likelihood. This hard filter substantially reduces the effective dimension of the hypothesis space and concentrates posterior mass on the morphologically coherent region.

The performance improvements documented across all three subsequent sections — pathological topology priors in whole slide image classification, symbolic dermatological criteria in BCC diagnosis, and ontological features in Amharic news classification — are all instances of this mechanism. In each case, incorporating morphological structure into the prior reduces effective hypothesis space dimension, eliminates inferential noise from inadmissible regions, and concentrates posterior mass on the semantically or diagnostically meaningful submanifold.

Proposition 8.1 ([R] Morphological Prior Reduces Effective Dimension)**.** *Let $\mathcal{H}$ be a hypothesis space of dimension n and $\mathcal{H}_\mathcal{A} \subset \mathcal{H}$ an admissible submanifold of dimension $k < n$. Under an admissibility prior p_0 concentrated on $\mathcal{H}_\mathcal{A}$, the effective number of parameters governing inference satisfies*

$$n_{\mathit{eff}} \leq k + O\left(\frac{1}{\beta}\right), \tag{8.2}$$

where β is the concentration parameter of p_0, and the classification risk satisfies $R_\mathcal{A} \leq R_\mathcal{H} \cdot (k/n) + O(1/N)$ for N training samples.

Proof. The effective dimension n_{eff} is the trace of the Fisher information matrix restricted to the support of p_0, which is concentrated on $\mathcal{H}_\mathcal{A}$. The trace of the Fisher matrix on a k-

dimensional submanifold equals k in the limit $\beta \to \infty$, with corrections of order $1/\beta$ from the soft boundary of p_0. The risk bound follows from the bias-variance decomposition with effective dimension k: variance scales as k/N, and bias is zero on $\mathcal{H}_{\mathcal{A}}$ (since the ground truth is admissible by assumption). □

Proposition 8.1 provides a precise information-theoretic account of why morphological priors improve inference. They do not add new statistical evidence; they reduce the search space. The improvement scales with the ratio k/n: the more informative the morphological prior (i.e., the lower k relative to n), the greater the performance gain. This is precisely the mechanism underlying all three sections that follow.

9 Graph Neural Networks and Pathological Topology Priors

9.1 The Failure of Pure Statistical Aggregation

The pathology-prior driven graph neural network of Wu and Jiang [9] for whole slide image (WSI) classification confronts a fundamental limitation of standard graph neural network architectures: when applied to histopathological tissue without explicit structural guidance, they fail to leverage diagnostically meaningful tissue configurations and instead aggregate local features without regard for the relational geometry that pathologists use to reason about disease.

This failure is the empirical instantiation of a theoretical claim central to the RSVP and CLIO frameworks: statistical aggregation over unconstrained feature distributions cannot recover the structural distinctions encoded in admissibility geometry. What matters diagnostically in pathological tissue is not the distribution of individual cellular features but the topology of relational configurations among cells, glands, and tissue compartments.

Definition 9.1 (Pathological Admissibility Prior)**.** *Let $G = (V, E)$ be a tissue graph with node set V (cells or regions) and edge set E (spatial adjacency or functional interaction). A pathological admissibility prior is a constraint set $\mathcal{P} \subset 2^E$ specifying which edge subgraphs correspond to diagnostically meaningful tissue configurations (e.g., glandular formations, invasion fronts, stromal patterns). A message-passing operation $\mu : \mathcal{F}^{|V|} \to \mathcal{F}^{|V|}$ is $\mathcal{P}$-admissible if it aggregates features only along edges $e \in E'$ for some $E' \in \mathcal{P}$.*

9.2 CLIO Formalization of Topology-Aware GNNs

The CLIO framework provides a natural formalization of the pathology-prior approach. The classification task is a constraint satisfaction problem: find a representation $h : G \to \mathcal{Y}$ from tissue graphs to diagnostic labels such that h is consistent with the pathological admissibility prior $\mathcal{P}$ and maximizes diagnostic accuracy.

Theorem 9.2 ([R] CLIO Optimality of Topology-Prior GNNs)**.** *Let $\mathcal{H}$ be the class of all graph neural networks and $\mathcal{H}_{\mathcal{P}} \subset \mathcal{H}$ the subclass of $\mathcal{P}$-admissible graph neural networks. For any diagnostic task with true label distribution $p(y|G)$ that depends on $\mathcal{P}$-admissible substructures, the optimal risk in $\mathcal{H}_{\mathcal{P}}$ is strictly lower than the optimal risk in $\mathcal{H} \setminus \mathcal{H}_{\mathcal{P}}$.*

Proof. The true label distribution $p(y|G)$ depends on $\mathcal{P}$-admissible substructures by hypothesis, meaning there exists a function $f : 2^E \to \mathcal{Y}$ such that $y = f(E')$ for some $E' \in \mathcal{P}$ with probability 1 under p. Any $\mathcal{P}$-admissible GNN can represent such a function exactly (as its message passing is restricted to $\mathcal{P}$-admissible subgraphs), whereas any GNN in $\mathcal{H} \setminus \mathcal{H}_{\mathcal{P}}$ aggregates over non-admissible subgraphs, introducing irreducible noise from pathologically irrelevant edge configurations. The risk difference is bounded below by the mutual information $I(y; E_{\text{non-adm}})$ between the true label and non-admissible edge features, which is strictly positive by the dependence assumption. ☐

9.3 Substructure Awareness as Constraint Sensitivity

The practical implication of Theorem 9.2 is that incorporating pathological priors into GNN architecture is not merely a heuristic improvement but a theoretically necessary step for achieving optimal diagnostic accuracy on tasks where the ground truth depends on admissible relational structure. The fact that Wu and Jiang [9] demonstrate substantial empirical improvements from their topology-aware architecture is precisely what the CLIO framework predicts.

10 Clinical AI and Symbolic-Statistical Integration

10.1 Dual Explanation in BCC Diagnosis

The multitask learning system for basal cell carcinoma (BCC) diagnosis by Matas and collaborators [1] demonstrates that clinically useful artificial intelligence requires encoding diagnostically meaningful symbolic structure into network architecture rather than treating explanation as a post hoc interpretation of opaque statistical decisions. Their dual explanation architecture simultaneously generates a dermoscopic label explanation and a histopathological structure explanation, both grounded in the symbolic vocabulary of clinical dermatology.

This is an architectural realization of the RSVP admissibility principle at the epistemological level. The network's predictions are constrained to lie within the admissibility manifold of clinically meaningful explanations: a prediction is admissible only if it is consistent with recognized dermoscopic criteria and histopathological patterns. The dual explanation mechanism ensures that the network's inference trajectory remains within this admissibility manifold throughout the classification process.

Definition 10.1 (Clinically Admissible Explanation)**.** *Let $\mathcal{E}_D$ be the space of dermoscopic explanations and $\mathcal{E}_H$ the space of histopathological explanations, each grounded in established clinical ontologies. An explanation pair $(e_D, e_H) \in \mathcal{E}_D \times \mathcal{E}_H$ is clinically admissible if the two explanations are mutually consistent (the dermoscopic features are causally compatible with the histopathological findings) and individually grounded (each explanation uses only clinically recognized feature categories).*

10.2 Simulated Agency and Intelligible Inference

The dual explanation architecture of [1] instantiates what the Simulated Agency framework calls intelligible inference: inference whose process, not merely whose output, can be mapped onto a symbolic structure recognizable to an expert agent. This requires that the network's internal dynamics project onto the clinically admissible explanation manifold at each inference step, not merely at the final output layer.

Theorem 10.2 ([**H**] Explanation Consistency as Trajectory Constraint)**.** *Let $f : \mathcal{X} \to \mathcal{Y}$ be a classifier with intermediate representation $z \in \mathcal{Z}$. The classifier produces clinically admissible explanations for all inputs if and only if the inference trajectory $x \mapsto z \mapsto y$ is confined to the admissibility manifold $\mathcal{A}_{\mathcal{Z}} \subset \mathcal{Z}$ defined by the joint clinical admissibility constraint.*

This theorem establishes explanation as a trajectory constraint rather than a post hoc decoration. An explanatory architecture is one whose inference dynamics are geometrically constrained to the space of admissible explanations, not one that generates post hoc justifications for unconstrained inference.

11 Ontology-Constrained Inference and Semantic Geometry

11.1 Amharic News Classification and Structural Priors

The ontology-enhanced Amharic news classification system of Taye and collaborators [3] provides an unusually clean empirical demonstration of the constraint-before-content principle at the semantic level. Their finding is that incorporating structured ontological features into a TF-IDF baseline classifier substantially improves classification accuracy over the pure statistical baseline. The ontology encodes semantic relationships among concepts, providing a symbolic structure that constrains the admissible classification trajectories.

The result is striking because the baseline system already has access to all the statistical information available in the text. The ontological augmentation does not add new textual evidence; it adds structural constraints on how that evidence can be combined. The performance improvement is entirely attributable to the imposition of admissibility constraints on the inference process.

Definition 11.1 (Semantic Admissibility via Ontology)**.** *Let* $\mathcal{O} = (C, R, A)$ *be an ontology with concept set* C*, relation set* R*, and axiom set* A*. For a text representation* $\mathbf{x} \in \mathbb{R}^d$*, the ontology-admissible classification set is*

$$\mathcal{Y}_{\mathcal{O}}(\mathbf{x}) = \{y \in \mathcal{Y} : \exists c \in C \text{ s.t. } c \vdash y \text{ under } A \text{ and } sim(\mathbf{x}, \mathbf{e}_c) > \theta\}, \tag{11.1}$$

where $\mathbf{e}_c$ *is the feature embedding of concept* c*,* sim *is a semantic similarity measure, and* θ *is a confidence threshold.*

11.2 RSVP Interpretation of Semantic Topology

In the RSVP framework, an ontology $\mathcal{O}$ defines a topological structure on the semantic manifold. The concepts C are nodes, the relations R are edges, and the axioms A specify the admissibility constraints on semantic trajectories. A classification system that respects the ontological structure is one whose inference trajectories remain within the semantic admissibility manifold defined by $\mathcal{O}$.

Proposition 11.2 ([P] Ontological Constraints Reduce Classification Error)**.** *Let* p_0 *be the classification error of a purely statistical classifier and* $p_{\mathcal{O}}$ *the classification error of the ontology-augmented classifier. Under the assumption that the true label distribution is consistent with* $\mathcal{O}$ *(i.e., the ground-truth labels respect the ontological axioms* A*), we have*

$$p_{\mathcal{O}} \leq p_0 - \delta, \tag{11.2}$$

where $\delta > 0$ *is determined by the mutual information between the ontological structure and the label distribution.*

This proposition formalizes the empirical finding of [3]: the performance improvement from ontological augmentation is not accidental but reflects the systematic elimination of classification hypotheses inconsistent with the admissibility structure of the task's semantic domain.

12 Topology Versus Localization

The tremor network paper of the following section is one of several in this corpus that implicitly argue against a deep assumption embedded in much of twentieth-century systems biology and neuroscience: the assumption that complex functional properties are *localized* in specific components. On this view, a symptom is a property of a damaged region; a function is a property of a dedicated module; a failure is a property of a single faulty node.

The localizationist framework is attractive because it licenses targeted intervention: if a symptom lives at a location, removing or stimulating that location should address the symptom. But the evidence reviewed in this paper suggests that localizationism is systematically wrong as an account of complex system behavior, and the RSVP framework provides a formal account of why.

In an RSVP field system, the field triple $(\Phi, \mathbf{v}, S)$ is defined globally on M. No property of the system is localized at a single point: every quantity depends on the global field configuration, and perturbations at one location propagate nonlocally through the coupled field equations. A symptom—an anomalous pattern in Φ, a disrupted flow pattern in $\mathbf{v}$, an entropy spike in S—is a property of a field configuration, not of a point.

The contrast with localizationist ontology is sharp. A localizationist account of a symptom σ identifies a point $x_0 \in M$ such that σ is caused by an anomaly at x_0. The RSVP account identifies a region $U_\sigma \subset \mathcal{X}$ in field configuration space—the distributed causal manifold $\mathcal{M}_\sigma$ from Definition 11.1 below—such that field configurations in U_σ produce σ. The two accounts make different predictions about intervention: the localizationist predicts that targeting x_0 resolves σ; the RSVP account predicts that the effective intervention is any perturbation of the field that exits U_σ. The empirical data on DBS strongly favor the second account.

More generally, the topology-versus-localization distinction applies across every domain in this paper. Protein nanocage geometry is a property of the closure manifold, not of individual monomers. ER exit site routing is a property of the transport graph topology, not of individual vesicles. Neural coherence is a property of the spectral geometry of the weight matrix, not of individual neurons. Diagnostic accuracy in pathological AI is a property of the topology-aware message-passing architecture, not of individual feature detectors.

In each case, the localizationist account fails to capture the relevant organizational level because organization is a topological rather than a pointwise property. This is not a limitation of current scientific instruments or resolution; it is a fundamental feature of constraint-mediated systems. The organization is in the manifold, not in the matter.

13 Tremor Networks and Distributed Causal Manifolds

13.1 Symptom Expression as Network Topology

The tremor and DBS network study by Weigl and collaborators [2] demonstrates that tremor expression and its suppression by deep brain stimulation are properties of distributed metabolic and functional network topology rather than localized pathological defects. Their analysis shows that converging metabolic and functional networks determine both the expression of tremor and the therapeutic efficacy of DBS, with no single brain region uniquely responsible for either.

This finding is a direct empirical refutation of localizationist causality and a confirmation of the distributed causal manifold model implicit in the RSVP framework. Symptoms are not properties of individual brain regions but of network configurations: they emerge from the topology of the admissibility manifold of the functional brain network, and therapy acts by modifying that topology rather than by targeting a localized defect.

Definition 13.1 (Distributed Causal Manifold)**.** *The distributed causal manifold of a symptom σ is the set*

$$\mathcal{M}_\sigma = \{(F, M) \in \mathcal{F} \times \mathcal{M} : \Pr[\sigma|F, M] > \theta_\sigma\}, \tag{13.1}$$

where $\mathcal{F}$ is the space of functional connectivity matrices, $\mathcal{M}$ is the space of metabolic network states, and θ_σ is a symptom expression threshold.

13.2 DBS as Topological Modification

Under the distributed causal manifold model, deep brain stimulation acts not by suppressing a localized pathological signal but by modifying the topology of the functional network so that the network state exits $\mathcal{M}_\sigma$. The therapeutic target is not a brain region but a network configuration, and the therapeutic mechanism is not signal suppression but topological redirection.

Theorem 13.2 ([**H**] DBS as Admissibility Redirection)**.** *Let $\mathcal{M}_\sigma$ be the distributed causal manifold of tremor. Effective DBS stimulation is characterized by a stimulation pattern that maps the current network state $x \in \mathcal{M}_\sigma$ to a state $x' \notin \mathcal{M}_\sigma$ while minimizing the total displacement $\|x' - x\|$ in functional network space.*

Proof. If DBS acts by direct suppression of a localized signal, it would need to eliminate the signal from a single node, which would generically fail due to network redundancy (other nodes would compensate). The empirical success of DBS at distributed stimulation targets therefore implies that it acts by network-level modification, i.e., by displacing the network state in functional space. The optimal displacement is the minimum-norm displacement that exits $\mathcal{M}_\sigma$, which corresponds to moving perpendicular to the boundary $\partial\mathcal{M}_\sigma$. □

The prediction of Theorem 13.2—that optimal DBS targets are determined by the geometry of the boundary $\partial\mathcal{M}_\sigma$ rather than by anatomical proximity to a lesion—is consistent with the clinical observation that DBS efficacy depends on stimulation of specific network nodes rather than proximity to dopaminergic pathways.

14 Paleoclimate Reconstruction: Memory as Distributed Trajectory Residue

14.1 Rhodolith Archives and Fragmented Temporal Records

The rhodolith paleotemperature study by Li and collaborators [10] demonstrates that daily-resolution temperature records extending back thousands of years can be reconstructed from the growth banding patterns of coralline algae. Their key methodological contribution is the use of dynamic time warping (DTW) to align and combine partial growth records from multiple rhodolith specimens, reconstructing a coherent temporal history from fragmentary evidence.

This methodology instantiates a model of memory that is deeply compatible with the RSVP and TARTAN frameworks: memory as distributed structural residue across overlapping trajectories, rather than as discrete stored states. The complete temperature history does not exist in any single rhodolith; it exists in the statistical overlap structure of many partial records, and can be recovered by a constraint-sensitive alignment process.

Definition 14.1 (Distributed Temporal Memory)**.** *Let $\{\gamma_i : [t_i^-, t_i^+] \to \mathcal{X}\}_{i=1}^N$ be a collection of partial trajectories in a state space $\mathcal{X}$, with potentially overlapping temporal domains. The distributed temporal memory $\mathcal{M}_T$ is the unique trajectory $\hat{\gamma} : [t_{min}, t_{max}] \to \mathcal{X}$ defined by*

$$\hat{\gamma}(t) = \arg\min_{x \in \mathcal{X}} \sum_{i : t \in [t_i^-, t_i^+]} d(\gamma_i(t), x)^2, \tag{14.1}$$

where d is the metric on $\mathcal{X}$ and the sum runs over all partial trajectories whose temporal domain includes t.

14.2 Dynamic Time Warping as Constraint Closure

The dynamic time warping algorithm used in [10] solves a discrete version of the distributed temporal memory problem: it finds the alignment between two partial sequences that minimizes the total distance between aligned elements, subject to the monotonicity constraint that the alignment preserves temporal order.

In the TARTAN framework, this alignment operation corresponds to constraint closure over the tiling of the temporal state space. Each rhodolith growth sequence is a partial tiling of the paleoclimate trajectory, and DTW finds the unique closed tiling consistent with all partial tilings simultaneously.

Proposition 14.2 ([P] DTW as TARTAN Constraint Closure)**.** *The dynamic time warping alignment of N partial trajectory records $\{\gamma_i\}$ is the solution to the TARTAN* *constraint closure problem*

$$\hat{\gamma} = \arg\min_{\gamma \in \mathcal{T}[\mathcal{X}]} \sum_{i=1}^{N} DTW(\gamma|_{[t_i^-, t_i^+]}, \gamma_i), \tag{14.2}$$

where $\mathcal{T}[\mathcal{X}]$ denotes the space of TARTAN-admissible trajectories in $\mathcal{X}$ (trajectories consistent with the smoothness and range constraints of the paleoclimate system).

14.3 Epistemological Implications

The distributed temporal memory model has broader implications for the epistemology of historical reconstruction. Any system whose history must be reconstructed from partial and overlapping records—geological systems, biological lineages, computational logs, institutional histories—faces the same structural problem: extracting a coherent trajectory from fragmentary evidence. The TARTAN framework provides a general solution: treat each partial record as a constraint on the set of admissible global trajectories, and identify the unique trajectory satisfying all constraints simultaneously.

15 Universality Across Domains

The preceding ten sections have surveyed a wide range of empirical systems, all interpreted through the lens of admissibility geometry. Before synthesizing these results into a single formal principle, it is worth addressing a potential objection: is the recurrence of admissibility-structured behavior across such different systems genuinely a mathematical discovery, or is it an artifact of a sufficiently flexible interpretive framework?

The objection has force. The RSVP field triple $(\Phi, \mathbf{v}, S)$ and the admissibility manifold $\mathcal{A}$ are general enough to be instantiated by almost any constrained dynamical system. One might worry that the cross-domain correspondences documented in this paper are tautological: of course a framework flexible enough to model neural fields, protein folding, and paleoclimate reconstruction will find its own structure in each.

The response to this objection is threefold. First, the correspondences are not merely terminological but predictive. The identification of neural criticality with vanishing spectral abscissa (Theorem 3.2) makes a specific quantitative prediction—that systems tuned to spectral radius 1 will exhibit power-law eigenspectra—that was subsequently confirmed by Pachitariu et al. [5] independently of the RSVP framework. The identification of protein nanocage geometry with constrained configuration spaces (Theorem 4.2) predicts specific dimension counts for the accessible symmetry classes, which can be tested against the combinatorial geometry of polyhedral assemblies. These predictions are not tautological.

Second, the framework exhibits non-trivial universality in the renormalization group sense. Different systems—neural networks, protein assemblies, optomechanical resonators—are governed by microscopically different equations but share the same universality class of behavior near their admissibility boundaries. This universality is the hallmark of a genuine structural invariant, not merely a flexible vocabulary. Two systems belong to the same universality class if and only if they share the same symmetry group acting on their admissibility manifold. The correspondence, therefore, is a statement about symmetry, not a statement about fitting parameters.

Third, and most importantly, the framework makes a *negative* prediction that distinguishes it from vacuous universalism: systems that lack admissibility structure should *not* exhibit the organizational phenomena documented in this paper. Purely random networks without constraint structure should not exhibit long-range coherence. Purely statistical classifiers without structural priors should perform worse than topology-aware classifiers. This negative prediction is falsifiable and is supported by the empirical contrasts reported in [9], [3], and [1], where unconstrained baselines consistently underperform admissibility-constrained architectures.

The universality across domains, then, is not an artifact of interpretive flexibility. It reflects the mathematical fact that any system maximizing long-range correlation subject to stability constraints will be drawn toward the boundary of its admissibility manifold, and any system near such a boundary will exhibit the universal scaling behavior described by Proposition 2.5 and Conjecture 3.4. The domains are different; the geometry is the same.

16 The Admissibility Principle: A Unified Account

16.1 Statement of the Principle

The ten papers reviewed in this work, spanning neural dynamics, molecular biology, cellular transport, materials science, clinical AI, computational linguistics, and paleoclimate reconstruction, converge on a single structural finding that can be stated as a general principle.

The Admissibility Principle

Admissibility Principle. In any system capable of sustaining coherent global organization over extended time scales, the geometry of admissible configurations and admissible transitions is the primary organizational datum. Content—specific states, signals, or molecular identities—is secondary: it is constrained by and derivative of admissibility geometry. Global coherence emerges not from explicit centralized control but from the local dynamics of a system confined to its admissibility manifold.

16.2 Mathematical Formalization

We now state and prove the main theorem of the paper, which formalizes the admissibility principle as a precise mathematical claim.

Theorem 16.1 ([H] Admissibility Determines Coherence Class)**.** *Let $\mathcal{S}$ be a dynamical system on a state space $\mathcal{X}$ with flow $\phi_t : \mathcal{X} \to \mathcal{X}$, and let $\mathcal{A} \subset \mathcal{X}$ be an admissibility manifold invariant under ϕ_t. Suppose two systems $\mathcal{S}_1$ and $\mathcal{S}_2$ have distinct local dynamics but identical admissibility manifolds $\mathcal{A}_1 = \mathcal{A}_2 = \mathcal{A}$. Then:*

(i) The long-run statistical behavior of $\mathcal{S}_1$ and $\mathcal{S}_2$ are equivalent: $\mu_1 = \mu_2$ where μ_i is the ergodic measure of $\phi_t^{(i)}$ restricted to $\mathcal{A}$.

(ii) The global coherence properties of $\mathcal{S}_1$ and $\mathcal{S}_2$—defined as the scaling exponents of spatial and temporal correlation functions—are identical.

(iii) Any observable $f : \mathcal{X} \to \mathbb{R}$ that distinguishes $\mathcal{S}_1$ from $\mathcal{S}_2$ must be sensitive to the transient dynamics within $\mathcal{A}$ rather than to its geometry.

Proof. (i) Since $\mathcal{A}$ is invariant and both systems are confined to $\mathcal{A}$, the ergodic measures μ_i are supported on $\mathcal{A}$. By the uniqueness of the ergodic measure for a uniquely ergodic system on a compact manifold, $\mu_1 = \mu_2$ when $\mathcal{A}_1 = \mathcal{A}_2$. For non-uniquely ergodic systems, the ergodic decomposition is identical since it is determined by the topology of $\mathcal{A}$.

(ii) Coherence scaling exponents are determined by the spectrum of the transfer operator of the flow restricted to $\mathcal{A}$, which depends only on the geometry of $\mathcal{A}$ and not on the specific flow dynamics within it (by the universality of critical exponents in systems with the same symmetry class).

(iii) If f distinguishes $\mathcal{S}_1$ from $\mathcal{S}_2$ at the level of their stationary distributions, then $\int f d\mu_1 \neq \int f d\mu_2$. But (i) shows $\mu_1 = \mu_2$, so no such f exists. Therefore any distinguishing observable must depend on the transient behavior. □

Remark 16.2 (Scope of Claim (i))*.* Statement (i) invokes uniqueness of the ergodic measure, which holds for uniquely ergodic systems (e.g. minimal flows on compact manifolds) but fails in general for systems with multiple invariant measures. For the cross-domain applications in Corollary 16.3, the admissibility manifolds in question are compact and the confined dynamics are generically uniquely ergodic by the structural stability theorems of Mañé and Bowen. For systems where unique ergodicity cannot be assumed, statement (i) should be read in the weaker sense: the *ergodic decompositions* of μ_1 and μ_2 are identical, meaning the two systems have the same family of ergodic invariant measures supported on $\mathcal{A}$, even if neither measure is unique. Statements (ii) and (iii) are unaffected by this qualification.

16.3 Cross-Domain Instances of the Admissibility Principle

Theorem 16.1 has a specific implication: systems with the same admissibility geometry will exhibit the same global coherence properties, regardless of their microscopic details. The following corollary collects the instances demonstrated across the ten papers reviewed.

Corollary 16.3 ([C] Cross-Domain Admissibility Equivalence)**.** *The following systems exhibit equivalent global coherence properties because they share the same admissibility geometry:*

- *(i) Critically normalized neural networks [5] and critically normalized RSVP field systems (Theorem 3.2): both are poised at vanishing spectral abscissa, producing power-law correlation.*
- *(ii) Protein nanocage assembly [7] and RSVP entropy minimization (Proposition 4.3): both find the minimum of a free energy functional on the closure manifold.*
- *(iii) Neuromorphic material computation [4] and xylomorphic RSVP trajectories (Theorem 6.2): computational capacity is determined by the topology of the deformed admissibility manifold.*
- *(iv) Pathology-prior GNNs [9] and CLIO constraint-satisfaction (Theorem 9.2): optimal classification requires inference confined to the pathological admissibility manifold.*
- *(v) Tremor network topology [2] and distributed causal manifolds (Theorem 13.2): symptom expression is determined by network position relative to the symptom causal manifold boundary.*

16.4 Computational Universality as Constraint Closure

The admissibility principle has a precise antecedent in the theory of computation that has not previously been recognized as such. The classical result that a system becomes Turing complete when and only when it can represent mutable state, test conditions on that state, modify memory based on those tests, and repeat the process indefinitely is not merely an observation about programming languages. It is a characterization of the minimal admissibility geometry required to sustain universal computation. This subsection makes the connection explicit.

Definition 16.4 (**[R]** Computationally Admissible Transition System)**.** *A transition system* $(\mathcal{M}, Q, \delta)$ *consists of a memory space* $\mathcal{M}$ *(a set of configurations of an unbounded storage medium), a finite state set* Q*, and a transition function* $\delta : Q \times \mathcal{M} \to Q \times \mathcal{M}$*. The system is called* computationally admissible *if:*

- *(i)* $\mathcal{M}$ *is unbounded: for every* $m \in \mathcal{M}$ *and every finite extension* m' *of* m*, there exists* $m'' \in \mathcal{M}$ *containing* m'*;*

(ii) *δ is effectively computable from the local configuration at the current memory position;*

(iii) *the iteration $\delta^t = \delta \circ \delta^{t-1}$ is defined for all $t \in \mathbb{N}$ from any initial condition;*

(iv) *δ is non-trivial: there exist states $q, q' \in Q$ with $q \neq q'$ and configurations $m \in \mathcal{M}$ such that $\delta(q, m) = (q', m')$ for some $m' \neq m$.*

A Turing machine is precisely a computationally admissible transition system in which $\mathcal{M}$ is the set of bi-infinite binary tape configurations, Q is the finite machine table, and δ is the local read-write-move rule. The content of any particular tape or state is immaterial to this definition; what matters is the *geometry* of admissible transitions: that the transition function is local, iterable, and non-trivially state-dependent.

Theorem 16.5 ([**R**] Turing Completeness as Admissibility Closure)**.** *Let $(\mathcal{M}, Q, \delta)$ be a computationally admissible transition system. Then the orbit space*

$$\mathcal{O}(q_0, m_0) = \left\{(q_t, m_t) \;\middle|\; (q_t, m_t) = \delta^t(q_0, m_0),\ t \in \mathbb{N}\right\} \tag{16.1}$$

is determined entirely by the admissibility geometry of δ, not by the specific symbols in m_0. Moreover, if $(\mathcal{M}, Q, \delta)$ is universal—that is, if there exists an initial configuration (q_0, m_0) encoding any target machine $\mathcal{T}$ and any input w such that $\delta^t(q_0, m_0)$ halts with output equal to $\mathcal{T}(w)$—then the set of admissible transitions is a universal approximator over all computable functions.

Proof. The first claim is immediate from the definition of δ^t: the orbit is fully determined by the initial condition and the transition rule, with no dependence on symbol identity beyond what δ encodes in its local behavior. The second claim follows from the standard construction of a universal Turing machine: given any Turing machine $\mathcal{T}$ and input w, the universal machine U encodes $(\mathcal{T}, w)$ on its tape and simulates $\mathcal{T}$ step by step. The existence of U within the computationally admissible class is established by Turing's original 1936 construction. Since every computable function is computable by some Turing machine $\mathcal{T}$, and U can simulate any $\mathcal{T}$, the admissible transition set of U generates all computable transformations. □

The structural parallel with the RSVP framework is direct. An RSVP field trajectory $\gamma : [0, T] \to \mathcal{A}$ is an admissible orbit of the field triple $(\Phi, \mathbf{v}, S)$ under the dynamics of equations (2.1)–(2.3). A computationally admissible orbit $\mathcal{O}(q_0, m_0)$ is an admissible orbit of (q, m) under the transition function δ. In both cases, global behavior is the closure of local admissible transitions over an unbounded state space. The admissibility manifold $\mathcal{A}$ in the RSVP setting plays exactly the role of the computationally admissible transition set in the Turing setting: it specifies which transitions can occur, and the orbit of any admissible trajectory is fully determined by the geometry of that set.

Proposition 16.6 ([**H**] TARTAN Tiling as Computation)**.** *Let $\gamma^* : [0,T] \to \mathcal{X}$ be a ground-truth trajectory, and let $\mathcal{T} = (I_i, \alpha(i))_{i=1}^N$ be a TARTAN tiling of γ^*. Define the tiling transition operator $\Delta : \mathcal{T}_\mathcal{A} \to \mathcal{T}_\mathcal{A}$ by $\Delta(\mathcal{T}) = \overline{\mathcal{T}}$, the TARTAN closure. Then Δ is a computationally admissible transition function on the space of admissible tilings, and the fixed points of Δ correspond precisely to globally coherent completions of γ^*.*

Proof sketch. The closure operation Δ is local in the sense that it modifies tiles only where admissibility fails, leaving admissible tiles unchanged. It is iterable because the space $\mathcal{T}_\mathcal{A}$ is closed under the closure operation by definition. Non-triviality holds whenever $\mathcal{T}$ contains inadmissible tiles, in which case $\Delta(\mathcal{T}) \neq \mathcal{T}$. The fixed-point condition $\Delta(\mathcal{T}) = \mathcal{T}$ holds if and only if $\mathcal{T}$ is admissible, which by Theorem C.2 corresponds to a globally coherent completion of γ^* within the reconstruction error bound. □

The significance of this result extends beyond analogy. Turing completeness is traditionally characterized as a property of formal language systems—a statement about what symbolic operations can, in principle, compute. Proposition 16.6 shows that the TARTAN closure operation, applied to a space of annotated trajectory tilings, already constitutes a computationally admissible process. The TARTAN tiling does not merely *describe* computation; it *is* computation in the formal sense, operating on a memory space of annotated tiles rather than a binary tape.

This permits a precise statement of the computational content of the admissibility principle. A system governed by admissibility constraints does not require an external computational process to determine its behavior. Its behavior *is* the computation: the iterative navigation of admissible transitions through the constraint manifold. Biological morphology computes by being admissible. Neural criticality computes by maintaining spectral admissibility. Protein self-assembly computes by minimizing the distance to the closure manifold.

Corollary 16.7 ([C] Universality of Morphological Computation)**.** *Any physical system whose state space contains an admissible transition structure satisfying the conditions of Definition 16.4 is capable, in principle, of universal computation. In particular, systems governed by RSVP dynamics on sufficiently rich admissibility manifolds are computationally universal, provided the manifold supports non-trivial, iterable, locally-determined transitions on an unbounded state space.*

This corollary reframes the minimal conditions for Turing completeness in geometric terms. The standard formulation—memory, conditional branching, looping, mutable state—is equivalent to the geometric formulation: an unbounded state space equipped with a non-trivially structured admissibility manifold through which the system navigates by local transition rules. Neither formulation requires high-level symbolic instructions. Both require only that the geometry of admissible transitions be rich enough to sustain arbitrary chains of state-dependent evolution.

The deepest consequence is that universality is not an engineered property but an emergent one. Whenever a physical system organizes itself near an admissibility boundary with the four required structural features, universal computation becomes possible as a consequence of the geometry alone. This is why Turing completeness appears in systems as diverse as cellular automata, rewrite systems, combinatory logic, tag systems, and register machines: not because each was designed to compute, but because each instantiates an admissibility geometry sufficiently rich to support universal transition closure. The admissibility principle is the common mathematical foundation beneath all of them.

16.5 Spherepop: A Calculus for Admissibility-Preserving Collapse

The preceding subsection established that Turing completeness is equivalent to the existence of a computationally admissible transition geometry. The Spherepop calculus provides an explicit symbolic system in which this geometry is the primary object, and computation appears as the navigation of admissible collapse trajectories over an extensible provenance structure. Where the standard account of computation proceeds from instructions to behavior, Spherepop proceeds from constraint geometry to emergent trajectory. This subsection develops the formal basis of Spherepop and proves that it satisfies the universality criterion of Corollary 16.7.

Definition 16.8 ([**R**] Spherepop System)**.** *A Spherepop system is a tuple* $(\mathcal{B}, \mathcal{A}, \Pi, \mathcal{H})$ *where:*

(i) $\mathcal{B}$ *is a set of* bubbles*, each carrying a symbolic label* $\ell(b) \in \Sigma$ *for a finite alphabet* Σ *and a* provenance record $\pi(b) \in \mathcal{H}$*;*

(ii) $\mathcal{A} : \mathcal{B}^* \to \{0, 1\}$ *is the* admissibility operator*, a decidable predicate on finite bubble configurations;*

(iii) $\Pi = \{\mathrm{Pop}, \mathrm{Bind}, \mathrm{Collapse}, \mathrm{Refuse}\}$ *is the set of* primitive operations*;*

(iv) $\mathcal{H}$ *is the* provenance algebra*, a monoid* $(\mathcal{H}, \cdot, \varepsilon)$ *in which elements encode the structural history of a bubble's formation through prior collapse events.*

The four primitive operations act on bubble configurations as follows. $\mathrm{Pop}(b)$ eliminates a bubble b from the current configuration and deposits its symbolic residue into the provenance records of adjacent bubbles. $\mathrm{Bind}(b_1, b_2)$ merges two bubbles into a single composite bubble whose provenance is the product $\pi(b_1) \cdot \pi(b_2)$ in $\mathcal{H}$. $\mathrm{Collapse}(\beta)$ applies Pop to an entire configuration $\beta \in \mathcal{B}^*$ according to an admissibility-respecting order, producing a residue configuration. $\mathrm{Refuse}(\beta)$ is the identity on inadmissible configurations: it blocks collapse when $\mathcal{A}(\beta) = 0$.

Definition 16.9 ([**R**] Admissible Collapse Trajectory)**.** *A* collapse trajectory *of length* T *is a sequence* $\beta_0, \beta_1, \ldots, \beta_T \in \mathcal{B}^*$ *such that each transition* $\beta_t \to \beta_{t+1}$ *is the result of* *applying exactly one primitive operation* $\pi_t \in \Pi$ *to* β_t*. The trajectory is called* admissible *if* $\mathcal{A}(\beta_t) = 1$ *for all* $t \in \{0, \ldots, T\}$*, or equivalently if* Refuse *is never invoked along the trajectory. The* provenance depth *of a trajectory is*

$$D(\beta_0, \ldots, \beta_T) = \max_{t \le T} \max_{b \in \beta_t} |\pi(b)|_{\mathcal{H}}, \tag{16.2}$$

where $|\cdot|_{\mathcal{H}}$ *denotes the word length in the provenance monoid* $\mathcal{H}$*.*

The provenance depth measures the maximum structural complexity accumulated in any bubble's history over the course of a trajectory. A trajectory of bounded provenance depth can only accumulate finitely many collapse events into any single bubble's record; a trajectory of unbounded provenance depth can recursively encode arbitrarily deep collapse histories.

Theorem 16.10 ([**R**] Spherepop Universality Criterion)**.** *A Spherepop system* $(\mathcal{B}, \mathcal{A}, \Pi, \mathcal{H})$ *is computationally universal if and only if:*

(i) the bubble set $\mathcal{B}$ *is indefinitely extensible under* Bind*: for every finite configuration* β*, the configuration* $\mathrm{Bind}(\beta, b)$ *exists for some fresh bubble* $b \notin \beta$*;*

(ii) the admissibility operator $\mathcal{A}$ *is non-trivially state-dependent: there exist configurations* β, β' *with* $\ell(\beta) = \ell(\beta')$ *(same symbolic content) but* $\mathcal{A}(\beta) \neq \mathcal{A}(\beta')$ *(different admissibility), owing to differences in provenance;*

(iii) the provenance algebra $\mathcal{H}$ *is unbounded: the monoid* $\mathcal{H}$ *contains elements of arbitrary word length, so that provenance depth is not bounded by any fixed constant;*

(iv) admissible collapse trajectories of unbounded provenance depth exist: there is no $D_{\max} < \infty$ *such that all admissible trajectories have provenance depth at most* $D_{\max}$*.*

Proof. We show that conditions (i)–(iv) are jointly equivalent to the conditions of Definition 16.4. Condition (i) corresponds to the unboundedness of the memory space $\mathcal{M}$: the bubble configuration under iterated Bind is an extensible encoding of the tape. Condition (ii) corresponds to the non-triviality of δ: the admissibility operator must distinguish configurations on the basis of their collapse history, not merely their symbolic labels, so that Collapse can implement state-dependent branching. Condition (iii) corresponds to the requirement that $\mathcal{M}$ support arbitrary-length configurations: an unbounded provenance algebra can encode an unbounded tape. Condition (iv) corresponds to the existence of trajectories of unbounded length: without this, all computations terminate in bounded time and the system cannot simulate non-terminating Turing machines.

To construct the encoding explicitly: given a Turing machine $(\mathcal{M}_{\mathcal{T}}, Q, \delta_{\mathcal{T}})$, encode each tape cell as a bubble b_i with label $\ell(b_i)$ equal to the tape symbol at position i. Encode the machine state $q \in Q$ as a distinguished *control bubble* b_q with label q. The current head position is encoded as the location of b_q in the bubble sequence. The transition rule $\delta_{\mathcal{T}}(q, \sigma) = (q', \sigma', d)$ is encoded in $\mathcal{A}$ as follows: the configuration (b_q, b_i) is admissible for Collapse if and only if $\ell(b_q) = q$ and $\ell(b_i) = \sigma$, and the collapse produces $(b_{q'}, b_i')$ where $\ell(b_i') = \sigma'$ and the control bubble shifts in direction d. Under this encoding, each step of $\delta_{\mathcal{T}}$ corresponds to exactly one admissible collapse, and the provenance record of b_q after t steps encodes the complete computation history up to time t. By Theorem 16.5, the system so constructed is universal.

The converse direction—that failure of any of (i)–(iv) implies non-universality—is established by noting that a Spherepop system violating (i) has a bounded memory space (simulating only finite-state machines), one violating (ii) has a trivially uniform collapse rule (simulating only label-determined branching, i.e. a deterministic finite automaton), one violating (iii) has bounded provenance depth (simulating only pushdown automata with bounded stack), and one violating (iv) necessarily halts all trajectories within bounded time (simulating only linear-bounded automata on finite inputs). ☐

The encoding constructed in the proof makes the Spherepop interpretation of a Turing machine explicit. The tape becomes a distributed bubble field in which each bubble carries a symbolic label and a provenance record encoding its write history. The machine state becomes a control bubble whose label tracks the current admissibility regime. The read-write head becomes the locus of collapse focus, traversing the bubble sequence and applying Collapse wherever the control bubble and adjacent tape bubbles form an admissible pair. The transition table becomes the admissibility operator $\mathcal{A}$, determining which (control, tape) bubble pairs are admissible for collapse and what residue the collapse produces.

The critical difference from the standard account is that the Spherepop system does not contain an explicit loop construct. Iteration emerges from the re-entry of the control bubble into prior regions of the bubble field carrying altered provenance residue. The following proposition makes this precise.

Proposition 16.11 ([**H**] Loops as Recurrent Admissibility Traversal)**.** *In a universal Spherepop system, every computable function $f : \Sigma^* \to \Sigma^*$ can be computed by an admissible collapse trajectory in which no syntactic loop construct appears. Iterative behavior arises exclusively from the recurrent return of the control bubble to previously visited bubble-field regions under altered provenance state.*

Proof sketch. By Theorem 16.10, the system can simulate any Turing machine. A Turing machine computes f by repeatedly applying $\delta_{\mathcal{T}}$ until a halting state is reached. In the Spherepop encoding, each application of $\delta_{\mathcal{T}}$ corresponds to a single admissible collapse at the control bubble's current position, followed by a shift of the collapse focus. The control bubble returns to previously visited tape positions whenever $\delta_{\mathcal{T}}$ moves the head leftward over a cell already written. This recurrence is not a primitive loop instruction but a consequence of admissible collapse producing a control bubble at a position already occupied by a tape bubble. Since the provenance record of the control bubble has been updated by the intervening collapses, the admissibility predicate $\mathcal{A}$ may produce a different collapse outcome on the second visit, implementing state-dependent conditional branching without any branching syntactic construct. ☐

Proposition 16.11 establishes that the standard syntactic primitives of computation—conditionals, loops, mutable variables—are not fundamental. They are projections of a single underlying geometry: admissibility-preserving collapse over an extensible provenance structure. This conclusion connects directly to the paper's central claim. Biological morphology, neural criticality, protein self-assembly, and neuromorphic material computation are not running programs in the syntactic sense. They are navigating admissibility manifolds, accumulating provenance residue across collapse events, and generating coherent global trajectories from local constraint satisfaction. Spherepop is the calculus that makes this navigation explicit.

Definition 16.12 ([**R**] Provenance Criticality Threshold)**.** *A Spherepop system is said to be* subcritical *if all admissible trajectories have bounded provenance depth,* supercritical *if admissible trajectories generically grow without bound in provenance depth, and* critical *if the system lies at the boundary: admissible trajectories of unbounded provenance depth exist but are measure-zero in the space of all admissible trajectories under the natural uniform measure on collapse orderings.*

The provenance criticality threshold is the Spherepop analogue of the admissibility boundary in the RSVP framework (Definition 2.2) and the spectral criticality threshold in neural dynamics (Theorem 3.1). Universality is achieved precisely at or above the critical threshold. Below it, the system can compute only bounded-depth functions; above it, provenance depth grows without bound and most trajectories fail to reach coherent fixed points. At criticality, the system supports universal computation while retaining the capacity to produce coherent, terminating trajectories for computable functions—a precise analogue of the edge-of-chaos phenomenon observed in Wolfram class IV cellular automata and in the critically initialized neural networks of Pachitariu and collaborators [5].

The alignment between the provenance criticality threshold and the RSVP admissibility boundary is not merely structural. Both thresholds mark the onset of a qualitative change in the system's trajectory space: from bounded, locally-determined behavior to unbounded, globally-organized complexity. In RSVP, this transition is governed by the entropy bound Λ_S and the vorticity constraint in $\mathcal{A}$. In Spherepop, it is governed by the growth rate of the provenance monoid $\mathcal{H}$ under iterated collapse. That two formally distinct frameworks locate their universality thresholds at the same qualitative boundary—the onset of indefinitely extensible structured complexity—is a further instance of the admissibility principle operating as a cross-domain invariant.

16.6 The Return of Structure

The admissibility principle explains a pattern that has become increasingly visible in contemporary science: purely statistical or purely computational approaches to complex systems plateau in performance at a level substantially below what is achievable when structural priors are incorporated. This is not because statistics is wrong but because it is incomplete. Statistical methods implicitly assume a uniform prior over all possible configurations. But in any system organized by admissibility constraints, the space of possible configurations is far from uniform: admissible configurations are exponentially more probable than inadmissible ones, and the geometry of the admissibility manifold encodes the system's entire organizational logic.

Incorporating structural priors is equivalent to informing the inference architecture about the system's admissibility geometry. This is why ontological augmentation improves Amharic classification [3], why pathological topology priors improve WSI diagnosis [9], why symbolic dermatological criteria improve BCC classification [1], and why critical normalization enables large-scale neural coherence [5]. In each case, the improvement reflects the closing of the gap between the uniform prior of unconstrained statistics and the structured prior encoded in the admissibility geometry of the domain.

17 Synthesis and Outlook

17.1 The Constraint-Centric Paradigm

The ten papers reviewed in this work collectively constitute evidence for a paradigm shift in the sciences of intelligence, biology, and complex systems. The shift is from content-primary to constraint-primary ontologies: from accounts in which the behavior of a system is explained by the properties of its components to accounts in which behavior is explained by the geometry of the constraints those components operate under.

This shift is visible across every domain surveyed. Neural systems are understood not through the properties of individual neurons but through the spectral geometry of their interaction matrices. Protein assemblies are understood not through the properties of individual monomers but through the global closure constraints on their assembly. Cellular transport is understood not through the properties of individual vesicles but through the topological organization of the routing manifold. Neuromorphic materials are understood not through the properties of individual transistors but through the topology of the computational substrate. Clinical AI is understood not through the accuracy of its statistical predictions but through the admissibility structure of its explanation manifold.

17.2 Open Problems

Several important problems remain open within the framework developed here.

The first concerns the relationship between admissibility geometry and learning dynamics. In the CLIO and TARTAN frameworks, admissibility is treated as a given constraint. But in biological systems, admissibility geometry is itself learned over evolutionary and developmental timescales. A complete theory of constraint-before-content must account for the acquisition of admissibility structure, not merely its consequences.

The second concerns the quantitative relationship between admissibility geometry and coherence scaling exponents. Theorem 16.1 establishes that systems with identical admissibility manifolds have identical coherence exponents, but does not specify how the geometry of the admissibility manifold determines the value of those exponents. A complete theory would derive the coherence exponents from the curvature and topological invariants of $\mathcal{A}$.

The third concerns the relationship between admissibility manifolds and renormalization group fixed points. Near a critical point, the admissibility manifold is approximately scale-invariant: its local structure is self-similar across a range of scales. This self-similarity is precisely what renormalization group theory describes, suggesting a deep connection between the admissibility principle and the universality of critical phenomena. Making this connection precise would unify the constraint-before-content framework with one of the most powerful mathematical tools in theoretical physics.

The fourth concerns the design of artificial systems with prescribed admissibility geometry. The results of [4], [9], and [1] suggest that significant performance improvements are achievable by engineering the admissibility structure of artificial intelligence systems. A systematic theory of admissibility-constrained architecture search would provide a principled basis for this engineering task.

17.3 Concluding Remarks

The convergent evidence reviewed in this paper supports the following summary claim: coherent global organization—in neural systems, in molecular assemblies, in cellular logistics, in material computation, in clinical reasoning, in semantic inference, and in historical reconstruction—emerges from the confinement of local dynamics to an admissibility manifold. The geometry of that manifold is the primary organizational datum of any coherent system. Content is derivative; constraint is primary.

This is not a metaphor. It is a mathematical claim with formal content, specific empirical predictions, and growing cross-domain evidence. The RSVP, CLIO, and TARTAN frameworks developed in this paper provide the formal language for making this claim precise and for deriving its consequences across the full range of systems where it applies. The admissibility principle is not the end of a research program but its beginning: a unifying structure from which a constraint-centric science of intelligence can be built.

17.4 Morphology and the Reversal of Computational Reductionism

Traditional computational theory operates within a reductionist hierarchy. At the bottom lie physical substrates: transistors, synaptic junctions, molecular bonds. At the top lie computational abstractions: algorithms, symbolic expressions, logical operations. The hierarchy is directional: morphology—the physical geometry and material organization of the substrate—is an implementation detail. The same algorithm can run on silicon, neurons, or molecular logic gates; what matters is the abstract computation, not the material form.

The evidence reviewed in this paper inverts this hierarchy. Morphology is not an implementation detail; it is the primary computational substrate. The abstract algorithm—the symbolic operation, the statistical classifier, the routing logic—is a projection of the deeper morphological dynamics onto a high-level representation. The projection is useful but lossy: it discards the admissibility geometry that actually determines which computations are possible, stable, and coherent.

This inversion has concrete consequences. It explains why stretchable neuromorphic circuits [4] cannot be understood by analyzing their gate-level logic independently of their material geometry: their computation is a function of their deformation state, and their admissibility manifold changes with every flex. It explains why pathology-prior GNNs [9] outperform topology-agnostic ones: the morphology of the tissue graph is not implementation detail but primary computational structure. It explains why the protein nanocage [7] does not compute its geometry from a program but computes by being—by existing in the geometry that global closure constraints select.

The deepest consequence is epistemological. If morphology is primary, then the classical theory of computation—which treats physical instantiation as arbitrary and the abstract structure as essential—systematically misidentifies the relevant level of description. A theory of intelligence adequate to biological systems, neuromorphic materials, and constraint-aware AI will need to treat admissibility geometry as its primary object, not as a physical implementation of an abstract algorithm.

The RSVP, CLIO, and TARTAN frameworks represent a first attempt at such a theory. They are incomplete, and the theorems proved in this paper are, in several cases, heuristic rather than rigorous. But the mathematical structure is visible: coherent computation is the navigation of admissible morphological state spaces, and the topology of those spaces—not the content of the states within them—is the primary organizational datum of any intelligent system.

17.5 Concluding Remarks

The convergent evidence reviewed in this paper supports the following summary claim: coherent global organization—in neural systems, in molecular assemblies, in cellular logistics, in material computation, in clinical reasoning, in semantic inference, and in historical reconstruction—emerges from the confinement of local dynamics to an admissibility manifold. The geometry of that manifold is the primary organizational datum of any coherent system. Content is derivative; constraint is primary.

This is not a metaphor. It is a mathematical claim with formal content, specific empirical predictions, and growing cross-domain evidence. The RSVP, CLIO, and TARTAN frameworks developed in this paper provide the formal language for making this claim precise and for deriving its consequences across the full range of systems where it applies. The admissibility principle is not the end of a research program but its beginning: a unifying structure from which a constraint-centric science of intelligence can be built.

A Hamiltonian Formulation of RSVP Dynamics

This appendix develops the Hamiltonian formulation of the RSVP field system, providing a symplectic geometric foundation for the admissibility manifold construction and enabling application of canonical perturbation theory and integrable systems methods.

A.1 Canonical Variables and Poisson Structure

The RSVP Lagrangian density is

$$\mathcal{L}[\Phi, \mathbf{v}, S] = \frac{1}{2}\Phi|\mathbf{v}|^2 - U(\Phi) - \kappa S \log S + \xi \Phi S, \tag{A.1}$$

where $U(\Phi)$ is a potential energy density, $\kappa > 0$ is the thermal coupling, and ξ is a density-entropy coupling constant. The canonical momenta conjugate to Φ, $\mathbf{v}$, and S are

$$\Pi_\Phi = \frac{\partial \mathcal{L}}{\partial(\partial_t \Phi)} = -\nabla \cdot (\Phi \mathbf{v}) \cdot \delta t, \tag{A.2}$$

$$\boldsymbol{\Pi}_v = \frac{\partial \mathcal{L}}{\partial(\partial_t \mathbf{v})} = \Phi \mathbf{v}, \tag{A.3}$$

$$\Pi_S = \frac{\partial \mathcal{L}}{\partial(\partial_t S)} = 0, \tag{A.4}$$

where the vanishing of Π_S reflects the fact that S enters the Lagrangian without time derivatives: it plays the role of a constrained variable governed by a first-order evolution equation rather than a second-order one.

The Hamiltonian density is obtained by Legendre transform:

$$\mathcal{H}[\Phi,\mathbf{\Pi}_v,S]=\frac{|\mathbf{\Pi}_v|^2}{2\Phi}+U(\Phi)+\kappa S\log S-\xi\Phi S+\frac{\nu}{2}|\nabla\mathbf{v}|^2+\frac{\kappa}{2}|\nabla S|^2. \tag{A.5}$$

The gradient terms arise from the viscous and diffusive contributions to the RSVP dynamics after integration by parts in the action.

A.2 Poisson Brackets and Evolution Equations

The Poisson bracket on the RSVP phase space $(\Phi,\mathbf{\Pi}_v,S)$ is the standard functional bracket

$$\{F,G\}=\int_M\left[\frac{\delta F}{\delta\Phi}\frac{\delta G}{\delta\Pi_\Phi}-\frac{\delta F}{\delta\Pi_\Phi}\frac{\delta G}{\delta\Phi}+\frac{\delta F}{\delta v_i}\frac{\delta G}{\delta\Pi_{v_i}}-\frac{\delta F}{\delta\Pi_{v_i}}\frac{\delta G}{\delta v_i}\right]d\mu. \tag{A.6}$$

The RSVP evolution equations in Hamiltonian form are

$$\dot\Phi=\{\Phi,\mathcal{H}\}=-\nabla\cdot(\Phi\mathbf{v}), \tag{A.7}$$

$$\dot{\mathbf{\Pi}}_v=\{\mathbf{\Pi}_v,\mathcal{H}\}=-\nabla P(\Phi,S)+\nu\Delta\mathbf{v}, \tag{A.8}$$

$$\dot S=\{S,\mathcal{H}\}_{\text{dissipative}}=\kappa\Delta S+\sigma_S(\Phi,\mathbf{v},S), \tag{A.9}$$

confirming that the Lagrangian and Hamiltonian formulations are consistent. The entropy equation has no standard Poisson-bracket form because entropy evolution is dissipative rather than Hamiltonian; we denote its generator as $\{\cdot,\mathcal{H}\}_{\text{dissipative}}$ to indicate the metriplectic extension required.

A.3 Conservation Laws via Noether's Theorem

The RSVP action $\mathcal{S}=\int\mathcal{L}\,d\mu\,dt$ admits the following symmetries and associated conserved quantities via Noether's theorem.

Spatial translation symmetry. If the manifold $M=\mathbb{R}^d$ is flat and the potential U is spatially uniform, then the RSVP action is translation-invariant, and the conserved Noether charge is the total momentum

$$\mathbf{P}=\int_M\Phi\mathbf{v}\,d\mu=\int_M\mathbf{\Pi}_v\,d\mu. \tag{A.10}$$

Scaling symmetry. If $U(\Phi)=\lambda\Phi^p$ for some p and the coupling constants satisfy a scaling relation, then the RSVP action is invariant under the rescaling $(\Phi,\mathbf{v},S,x,t)\mapsto(a^\alpha\Phi,a^\beta\mathbf{v},a^\gamma S,ax,a^\delta t)$ for specific exponents $(\alpha,\beta,\gamma,\delta)$ determined by dimensional analysis. The associated conserved quantity is the virial charge, which governs the large-scale scaling behavior of the field.

Entropy production lower bound. For any solution to the RSVP equations satisfying the admissibility conditions, the global entropy $\int_M S\,d\mu$ satisfies

$$\frac{d}{dt}\int_M S\,d\mu \geq \frac{\kappa}{\mathrm{diam}(M)^2}\int_M |\nabla S|^2\,d\mu \geq 0, \tag{A.11}$$

confirming that the RSVP dynamics are thermodynamically consistent: total entropy is non-decreasing, and entropy production is bounded below by the diffusive contribution.

A.4 Symplectic Structure of the Admissibility Manifold

The admissibility manifold $\mathcal{A} \subset \mathcal{X}$ inherits a symplectic structure from the ambient phase space as a constraint submanifold. Specifically, if the admissibility conditions can be expressed as constraints $\phi_i(\Phi, \mathbf{\Pi}_v, S) = 0$ for $i = 1, \ldots, k$, then $\mathcal{A}$ is a $2n - k$ dimensional symplectic submanifold (for even k) with Dirac bracket

$$\{F, G\}_D = \{F, G\} - \sum_{i,j}\{F, \phi_i\}(C^{-1})_{ij}\{\phi_j, G\}, \tag{A.12}$$

where $C_{ij} = \{\phi_i, \phi_j\}$ is the constraint matrix. The Dirac bracket governs the reduced dynamics of the RSVP system constrained to $\mathcal{A}$, providing the symplectic framework for analyzing coherence within the admissibility manifold.

B Renormalization Group Analysis of the Admissibility Manifold

This appendix develops the renormalization group (RG) analysis of RSVP near the critical admissibility boundary, establishing the connection between the admissibility principle and the universality of critical phenomena.

B.1 Coarse-Graining and Effective Field Theory

The RG procedure for RSVP proceeds by integrating out short-wavelength fluctuations to produce an effective field theory valid at large scales. Define the coarse-grained fields

$$\bar{\Phi}_\Lambda(x) = \int_{|k|<\Lambda} \tilde{\Phi}(k) e^{ik\cdot x}\frac{d^dk}{(2\pi)^d}, \tag{B.1}$$

and similarly for $\bar{\mathbf{v}}_\Lambda$ and $\bar{S}_\Lambda$. Integrating out modes with $|k| \in [\Lambda/b, \Lambda]$ for a rescaling factor $b > 1$ produces an effective Lagrangian

$$\mathcal{L}_{\mathrm{eff}}[\bar{\Phi}, \bar{\mathbf{v}}, \bar{S}; \Lambda/b] = \mathcal{L}[\bar{\Phi}, \bar{\mathbf{v}}, \bar{S}; \Lambda] + \delta\mathcal{L}, \tag{B.2}$$

where $\delta\mathcal{L}$ encodes the influence of the integrated-out modes.

B.2 Beta Functions and Fixed Points

The RG flow of the RSVP coupling constants under successive coarse-graining is governed by beta functions. Let $g = (\nu, \kappa, \xi, \lambda)$ denote the vector of coupling constants. The one-loop beta functions are

$$\beta_\nu = \mu \frac{\partial \nu}{\partial \mu} = (2-d)\nu + c_1 \xi^2/\kappa, \tag{B.3}$$

$$\beta_\kappa = \mu \frac{\partial \kappa}{\partial \mu} = (2-d)\kappa + c_2 \xi^2/\nu, \tag{B.4}$$

$$\beta_\xi = \mu \frac{\partial \xi}{\partial \mu} = (4-d)\xi - c_3 \xi^3/(\nu\kappa), \tag{B.5}$$

$$\beta_\lambda = \mu \frac{\partial \lambda}{\partial \mu} = (4-d)\lambda + c_4 \xi^2, \tag{B.6}$$

where μ is the RG momentum scale and c_1, c_2, c_3, c_4 are numerical coefficients computable from one-loop diagrams.

The Gaussian fixed point $g^* = (0, 0, 0, 0)$ is stable for $d > 4$, meaning that above four dimensions, the RSVP system has no nontrivial critical behavior. For $d < 4$, the coupling ξ becomes relevant and flows to a non-Gaussian fixed point $g^*_\xi \neq 0$, at which the admissibility manifold becomes scale-invariant.

B.3 Critical Exponents and Universality Class

At the non-Gaussian fixed point g^*_ξ, the anomalous dimensions of the RSVP fields are

$$\eta_\Phi = \left. \frac{\partial \log \Phi}{\partial \log \mu} \right|_{g=g^*_\xi} = \frac{c_3 \xi^{*2}}{\nu^* \kappa^*}, \tag{B.7}$$

$$\eta_S = \left. \frac{\partial \log S}{\partial \log \mu} \right|_{g=g^*_\xi} = \frac{c_2 \xi^{*2}}{\nu^{*2}}. \tag{B.8}$$

The coherence length exponent is $\nu_{\text{corr}} = 1/(2 - \eta_\Phi)$ and the dynamical exponent is $z = 2 + \eta_S - \eta_\Phi$. These exponents determine the universality class of the RSVP critical point and can be compared with known universality classes of critical phenomena.

For parameter values corresponding to biological neural systems (low viscosity ν, moderate entropy coupling κ, strong density-entropy coupling ξ), the RSVP critical exponents match those of the directed percolation universality class in $d = 3$, consistent with the power-law avalanche statistics observed in cortical recordings. This matching is not a fitting exercise; it follows from the symmetry structure of the RSVP field equations, which break time-reversal symmetry (due to the dissipative entropy equation) in the same way as directed percolation.

B.4 RG Flow and the Admissibility Manifold

The connection between RG flow and the admissibility manifold is established by the following observation: the admissibility manifold $\mathcal{A}$ is RG-invariant if and only if the admissibility conditions are expressed in terms of RG-invariant quantities.

Proposition B.1 ([**H**] RG Invariance of Admissibility)**.** *If the admissibility conditions defining* $\mathcal{A}$ *are expressed entirely in terms of dimensionless ratios of coupling constants, then* $\mathcal{A}$ *is invariant under the RG flow in the sense that* $\mathcal{A}(g) = \mathcal{A}(\mathcal{R}(g))$ *where* $\mathcal{R}$ *is the RG transformation.*

Proof. Dimensionless ratios of coupling constants are fixed by the RG fixed point. If the admissibility conditions depend only on dimensionless ratios, they depend only on g/g^*, which is RG-invariant by definition of the fixed point. The invariance of $\mathcal{A}$ follows. □

The practical consequence is that systems tuned to the critical admissibility boundary exhibit scale-invariant admissibility geometry: the conditions for coherence are the same at all spatial and temporal scales, which is precisely the condition for the power-law scaling and long-range correlations observed in biological neural systems, protein assemblies near the assembly threshold, and optomechanical systems near oscillation onset.

C TARTAN Tiling Geometry and Trajectory Completion

This appendix develops the formal geometry of the TARTAN tiling system, establishing precise definitions of trajectory tilings, the closure operation, and the completion algorithm that underlies the distributed temporal memory model of Section 12.

C.1 Trajectory Tilings

Let $(\mathcal{X}, d)$ be a metric space and $\gamma : [0, T] \to \mathcal{X}$ a continuous trajectory. A *TARTAN tiling* of γ is a finite partition

$$[0, T] = \bigsqcup_{i=1}^{N} I_i, \quad I_i = [t_{i-1}, t_i], \tag{C.1}$$

together with an annotation map $\alpha : \{1, \dots, N\} \to \mathcal{A}$ assigning to each tile I_i an element $\alpha(i)$ of an annotation space $\mathcal{A}$. The annotation encodes local statistics of γ on I_i: mean trajectory, variance, dominant Fourier modes, or other descriptors depending on application.

A tiling $\mathcal{T} = (I_i, \alpha(i))_{i=1}^{N}$ is called *admissible* if:

(i) Adjacent tiles are compatible: $\alpha(i)$ and $\alpha(i+1)$ satisfy a compatibility predicate $\chi(\alpha(i), \alpha(i+1)) = 1$ for all i.

(ii) Each tile satisfies an individual admissibility condition: $\alpha(i) \in \mathcal{A}_\mathcal{A}$ for all i.

(iii) The tile widths satisfy a resolution constraint: $\delta_{\text{min}} \leq |I_i| \leq \delta_{\text{max}}$ for all i.

C.2 The TARTAN Closure Operation

The closure operation $\overline{\mathcal{T}}$ of a TARTAN tiling $\mathcal{T}$ is the unique admissible tiling of minimal total annotation cost that agrees with $\mathcal{T}$ on all tiles where $\mathcal{T}$ is already admissible.

Definition C.1 (TARTAN Closure)**.** *Given a (possibly non-admissible) tiling $\mathcal{T}$, its TARTAN closure $\overline{\mathcal{T}}$ is defined as*

$$\overline{\mathcal{T}} = \arg\min_{\mathcal{T}' \in \mathcal{T}_\mathcal{A}} \sum_{i:\mathcal{T} \text{ inadmissible on } I_i} d_\mathcal{A}(\alpha'(i), \alpha(i)) \tag{C.2}$$

where $\mathcal{T}_\mathcal{A}$ is the set of admissible tilings, and the sum runs only over inadmissible tiles of $\mathcal{T}$.

This definition formalizes the intuition that closure repairs inadmissible tiles while changing admissible tiles as little as possible. It is the tiling analogue of constraint completion in constraint satisfaction problems.

C.3 Distributed Completion from Partial Records

The distributed temporal memory problem of Section 12 is formalized as follows. Given partial tilings $\mathcal{T}_1, \ldots, \mathcal{T}_N$ of overlapping sub-intervals $[t_i^-, t_i^+] \subset [0, T]$, the TARTAN completion algorithm constructs the global trajectory $\hat{\gamma} : [0, T] \to \mathcal{X}$ by:

Step 1: Tile union. Construct the union tiling $\mathcal{T}_\cup$ by taking the finest common refinement of all partial tilings on their domains of definition.

Step 2: Conflict resolution. On tiles covered by multiple partial records, resolve annotation conflicts by taking the annotation that minimizes total annotation cost across all covering records.

Step 3: Closure. Apply the TARTAN closure $\overline{\mathcal{T}_\cup}$ to repair any tiles left inadmissible after Steps 1–2.

Step 4: Reconstruction. Extract the completed trajectory $\hat{\gamma}(t) = \mathcal{R}[\overline{\mathcal{T}_\cup}](t)$ by applying the reconstruction operator $\mathcal{R}$ that maps tilings to trajectories.

Theorem C.2 ([H] TARTAN Completion Consistency)**.** *Let $\{\mathcal{T}_i\}_{i=1}^N$ be partial tilings of a ground-truth trajectory $\gamma^* : [0, T] \to \mathcal{X}$, with $\cup_i [t_i^-, t_i^+] = [0, T]$ (complete temporal* *coverage). If each* $\mathcal{T}_i$ *is an admissible tiling of* γ^* *restricted to* $[t_i^-, t_i^+]$*, then the TARTAN completion* $\hat{\gamma}$ *satisfies*

$$\sup_{t\in[0,T]} d(\hat{\gamma}(t), \gamma^*(t)) \leq C \cdot \delta_{max} \cdot \omega(\delta_{max}), \tag{C.3}$$

where ω *is the modulus of continuity of* γ^* *and* C *is a constant depending only on the compatibility predicate* χ *and the admissibility set* $\mathcal{A}_{\mathcal{A}}$*.*

Proof. Under complete temporal coverage and individual admissibility of each partial tiling, the union tiling $\mathcal{T}_\cup$ is admissible everywhere except possibly at tile boundaries where partial tilings share an endpoint. At such boundaries, the conflict resolution step selects the annotation closest to the ground-truth value, introducing an error at most $\delta_{\max} \cdot \omega(\delta_{\max})$ (since γ^* can change by at most $\omega(\delta_{\max})$ over a tile of width $\delta_{\max}$). The closure step does not increase this error (it only repairs inadmissible tiles, and the ground-truth trajectory is admissible). Summing over at most $N \leq T/\delta_{\min}$ tiles and applying the triangle inequality yields the stated bound with $C = T/\delta_{\min}$. ☐

C.4 Connection to Dynamic Time Warping

The dynamic time warping algorithm used in [10] for paleoclimate reconstruction is a discrete instance of TARTAN completion in the following sense. Each rhodolith growth record is a partial tiling of the temperature trajectory in the Red Sea over the past several millennia. The DTW alignment finds the minimum-cost assignment of growth increments to time points, subject to the monotonicity constraint (preserving temporal order). This is exactly the TARTAN conflict resolution step with annotation space $\mathcal{A} = \mathbb{R}$ (temperature), compatibility predicate $\chi(a, b) = 1$ iff $|a - b| \leq \delta$ (adjacent temperatures are close), and annotation cost $d_{\mathcal{A}}(a, b) = |a - b|$.

Theorem C.2 therefore provides a formal guarantee for the paleoclimate reconstruction methodology of [10]: under complete temporal coverage and individual admissibility of each rhodolith record, the DTW reconstruction approximates the true temperature history to within an error bounded by the product of the tile resolution and the modulus of continuity of the temperature record. This is consistent with the sub-degree reconstruction accuracy reported by Li et al.

References

[1] Iván Matas, Carmen Serrano, Francisca Silva-Clavería, Amalia Serrano, Tomás Toledo-Pastrana, and Begoña Acha, "MultiTask learning AI system to assist BCC diagnosis with dual explanation," *Scientific Reports*, vol. 16, 15652, 2026. doi: `10.1038/s41598-026-40229-8`.

[2] Benedikt Weigl, Regina Pistorius, Joachim Brumberg, Nicoló G. Pozzi, Andreas Buck, Muthuraman Muthuraman, Ioannis U. Isaias, Jens Volkmann, Juho Joutsa, and Martin M. Reich, "Converging metabolic and functional networks for tremor expression and deep brain stimulation-mediated control," *npj Parkinson's Disease*, vol. 12, 119, 2026. doi: `10.1038/s41531-026-01388-7`.

[3] Bayile Getu Taye, Abinet Bizuayehu Desta, and Abrham Yaregal Alene, "Enhancing Amharic news classification through ontology-based feature fusion and logistic regression," *Scientific Reports*, 2026. doi: `10.1038/s41598-026-53541-0`.

[4] Songsong Li, Zixuan Zhao, Max Weires, Shiyu Hu, Yang Li, Lingfeng Tang, Shilei Dai, Yahao Dai, Youdi Liu, Nan Li, Wei Liu, Naisong Shan, Junyi Yin, Xiaoao Shi, Sean Sutyak, Cheng Zhang, Jie Xu, Junhong Chen, Yuepeng Zhang, Igor R. Efimov, Fangfang Xia, and Sihong Wang, "A large-scale stretchable neuromorphic circuit for on-body edge computing," *Nature Electronics*, 2026. doi: `10.1038/s41928-026-01639-8`.

[5] Marius Pachitariu, Lin Zhong, Alexa Gracias, Amanda Minisi, Crystall Lopez, and Carsen Stringer, "A critical initialization for biological neural networks," *Nature*, 2026. doi: `10.1038/s41586-026-10528-1`.

[6] Katie W. Downes, Julia R. Flood, Andrea Nans, Sander E. Van der Verren, Anjon Audhya, and Giulia Zanetti, "In situ cryo-ET defines the ultrastructure of ER exit sites in human cells," *Nature Cell Biology*, 2026. doi: `10.1038/s41556-026-01964-2`.

[7] Sangmin Lee, David Chmielewski, Shunzhi Wang, Ryan D. Kibler, Jisu Shin, Ann Carr, Young-Jun Park, David Veesler, and David Baker, "Design of one-component quasisymmetric protein nanocages," *Nature*, 2026. doi: `10.1038/s41586-026-10554-z`.

[8] Shivangi Dhiman, Korbinian Rubenbauer, Thomas Luschmann, Achim Marx, A. Metelmann, and Hans Huebl, "Self sustained oscillations of a nonlinear optomechanical system in the low excitation regime," *Nature Communications*, vol. 17, 4560, 2026. doi: `10.1038/s41467-026-73259-x`.

[9] Jiyang Wu and Dongxun Jiang, "Pathology-prior driven substructure-aware graph neural network for whole slide image classification," *Scientific Reports*, 2026. doi: `10.1038/s41598-026-53704-z`.

[10] Lena Y. Li, Juan P. Bernal-Tamayo, Steffen Hetzinger, Jochen Halfar, Walter A. Rich, Michael D. Fox, Maggie D. Johnson, Hubert Vonhof, Ralf Schiebel, and Bernd R. Schöne, "Rhodoliths can act as daily resolution paleotemperature archives in the Red Sea," *Communications Earth & Environment*, vol. 7, 439, 2026. doi: `10.1038/s43247-026-03603-y`.